

# ADTSC Science Highlights 2008

## *Preface*

The Theory, Simulation, and Computation (TSC) Directorate is central to the huge national need for new generations of ideas, concepts, and methodologies to improve the fidelity, reliability, certainty and usability of tools to guide and interpret experiments, and provide prediction and control for complex phenomena and systems: the “Information Science and Technology” capability between data and prediction. Working with teams across the Laboratory, TSC undertakes major interdisciplinary challenges of integrating theory, modeling, simulation, high-performance computing, and visualization with experimental and other data to expand LANL’s predictive and uncertainty quantification capability. The core nuclear weapons program depends critically on the viability of this approach. However, virtually every emerging mission initiative at LANL also relies centrally on a similar integrated capability: threat reduction, energy security, manufacturing science, and many discipline frontiers from biology to materials science to cosmology.

Each year, TSC staff select a collection of brief topical reports on their recent research results. This annual collection is not intended to be complete. Rather it aims to demonstrate, through timely examples, the interdisciplinary energy and progress across many scientific and program frontiers. The collection chosen this year represents not only work for the National Nuclear Security Administration but also showcases the breadth of our “work with others” portfolio, including substantial investments from the DOE Office of Science, the National Institutes of Health, the National Science Foundation, and the Department of Homeland Security and other Threat Reduction agencies. This portfolio represents major outreach activities to universities, industry and other research laboratories, and is a basis for providing the best science and scientists to both current and future national security missions at LANL.

Alan R. Bishop  
*Associate Director, ADTSC*

ADTSC web site: <http://www.lanl.gov/orgs/adtsc>

[www.lanl.gov/orgs/adtsc/publications.php](http://www.lanl.gov/orgs/adtsc/publications.php)

## Associate Directorate for Theory, Simulation, and Computation (ADTSC)

---

**Alan R. Bishop**  
*Associate Director*

**Paul J. Dotson**  
*Deputy Associate Director*

**Andrew White**  
*Deputy Associate Director for  
High-Performance Computing*

**Stephen R. Lee**  
*CCS Division Leader*

**Antonio Redondo**  
*T Division Leader*

**Alyn Ford**  
*CTN Division Leader, Acting*

**John F. Morrison**  
*HPC Division Leader*

**Audrey Archuleta**  
*ADTSC Chief of Staff*

## Publications Team

---

**Susanne King,**  
**Kathy Pallis**  
*Editors*

**Sharon Mikkelsen**  
*Design*

**Guadalupe D. Archuleta**  
*Print Coordinator*





# Table of Contents



Preface .....	i
Funding Cross Reference .....	xi
Author Cross Reference .....	xiii
Organizational Abbreviation Legend .....	xvii

## Advanced Computational Architecture



<b>Direct Numerical Simulation of Turbulent Flows on Roadrunner .....</b>	<b>2</b>
Jamaludin Mohd-Yusof, Daniel Livescu, Mark R. Petersen, CCS-2; Nehal Desai, CCS-1	
<b>Fortran Development Tools: Providing a Roadmap for Application Development on Advanced Computer Architectures .....</b>	<b>4</b>
Craig E. Rasmussen, Christopher Rickett, CCS-1; Dan Quinlan, Lawrence Livermore National Laboratory; Matthew Sottile, Univ. of Oregon	
<b>Performance of Roadrunner under a Realistic Application Workload .....</b>	<b>6</b>
Kevin J. Barker, M. Kei Davis, Adolfo Hoisie, Darren J. Kerbyson, Mike Lang, Scott Pakin, Jose C. Sancho, CCS-1	
<b>Exploiting the Cell BE Architecture on Roadrunner for Large-scale Molecular Dynamics .....</b>	<b>8</b>
Sriram Swaminarayan, CCS-2; Timothy C. Germann, T-12; Kai Kadau, T-14	
<b>Petascale Plasma Physics Simulations on Roadrunner Using VPIC .....</b>	<b>10</b>
Benjamin Bergen, CCS-2; Brian J. Albright, Kevin J. Bowers, William Daughton, Thomas J. Kwan, Lin Yin, X-1 -PTA	
<b>Open MPI and the Roadrunner Project .....</b>	<b>12</b>
David Daniel, CCS-1; Ralph Castain, ISR-4; Brian Barrett, Sandia National Laboratories; Galen Shipman, Oak Ridge National Laboratory	
<b>Adaptation of Milagro Implicit Monte Carlo to the Roadrunner Hybrid Architecture .....</b>	<b>14</b>
Timothy M. Kelley, Paul J. Henning, CCS-2	
<b>Future and Emerging Computing Paradigms and Machines .....</b>	<b>16</b>
Christof Teuscher, Anders A. Hansson, CCS-3; Elshan A. Akhadov, MPA-CINT; Hsing-Lin Wang, C-PCS	
<b>Exploiting the Cell Architecture on Roadrunner for Long-time Atomistic Dynamics .....</b>	<b>18</b>
Danny Perez, Arthur F. Voter, T-12; Sriram Swaminarayan, CCS-2	

## Applied Mathematics, Fluid Dynamics, and Magnetohydrodynamics

<b>Simulations of Forced Isotropic Compressible Turbulence .....</b>	<b>22</b>
Mark R. Petersen, Daniel Livescu, Jamaludin Mohd-Yusof, CCS-2	
<b>Shock-turbulence Interaction .....</b>	<b>24</b>
Jamaludin Mohd-Yusof, Daniel Livescu, Mark R. Petersen, CCS-2	
<b>Convergence Effects on Acceleration-driven Instability .....</b>	<b>26</b>
Huidan Yu, Daniel Livescu, CCS-2	




<b>The Marginalization Paradox .....</b>	<b>28</b>
Timothy C. Wallstrom, T-13	
<b>The Importance of Fluctuations in Fluid Mixing .....</b>	<b>30</b>
Kai Kadau, John L. Barber, Timothy C. Germann, T-14; Pierre Carlès, Université Paris; Berni J. Alder, Lawrence Livermore National Laboratory; Charles Rosenblatt, Zhibin Huang, Case Western Reserve Univ.	
<b>The Scaling Properties of Atomistic Fluid Dynamics Simulations .....</b>	<b>32</b>
John L. Barber, Kai Kadau, T-14; Timothy C. Germann, T-12; Berni J. Alder, Lawrence Livermore National Laboratory	
<b>Scalable Solvers for 3D Extended Magnetohydrodynamics .....</b>	<b>34</b>
Luis Chacón, T-15; Bobby Philip, T-7	
<b>An Optimal Robust Equidistribution Method for 2D Grid Generation Based on Monge-Kantorovich Optimization .....</b>	<b>36</b>
Gian L. Delzanno, Luis Chacón, John M. Finn, Yeo-Jin Chung, Giovanni Lapenta, T-15	
<b>High-order Mimetic Finite Difference Methods on Arbitrary Meshes .....</b>	<b>38</b>
Konstantin Lipnikov, T-7; Vitaliy Gyrya, Pennsylvania State Univ.	
<b>A Multilevel Multiscale Mimetic (M3) Method for Two-phase Flows in Porous Media.....</b>	<b>40</b>
Konstantin Lipnikov, David Moulton, Daniil Svyatskiy, T-7	
<b>New Monotone Finite Volume Scheme for Anisotropic Diffusion .....</b>	<b>42</b>
Konstantin Lipnikov, Mikhail Shashkov, Daniil Svyatskiy, T-7; Yuri Vassilevski, Russian Academy of Sciences	
<b>Predicting the Small Scales of <math>\alpha</math>-Models for Turbulence .....</b>	<b>44</b>
Susan Kurien, T-7; Evelyn Lunasin, Univ. of California, Irvine	
<b>Understanding the Small Scales in Rotating Stratified Flow .....</b>	<b>46</b>
Susan Kurien, T-7; Beth Wingate, CCS-2; Leslie Smith, Univ. of Wisconsin	
<b>Modeling Multiscale-multiphase-multicomponent Subsurface Reactive Flows Using Advanced Computing.....</b>	<b>48</b>
Peter Lichtner, Chuan Lu, EES-6; David Moulton, Bobby Philip, T-7	
<b>Astrophysics and Cosmology</b>	
<b>Studying Gamma-ray Burst Progenitors under the NASA Swift Guest Investigator Program .....</b>	<b>52</b>
Chris L. Fryer, Gabriel Rockefeller, Patrick A. Young, CCS-2	
<b>Astronomical Data Analysis with Commodity Components .....</b>	<b>54</b>
Michael S. Warren, T-6; John Wofford, Columbia Univ.	
<b>Going Nonlinear with Dark Energy .....</b>	<b>56</b>
Salman Habib, T-8; Michael S. Warren, Darren Reed, T-6; Katrin Heitmann, ISR-1; Martin White, Lawrence Berkeley National Laboratory	
<b>Supernovae, Gamma-ray Bursts, and Nucleosynthesis .....</b>	<b>58</b>
Alexander Heger, T-6	




## Atomic, Nuclear and High Energy Physics

<b>Bose-Einstein Condensation in the Relativistic Ideal Bose Gas .....</b>	<b>62</b>
George A. Baker, Jr., T-11	
<b>Microscopic Calculation of the Pre-equilibrium Emission for Neutron Scattering on Actinides .....</b>	<b>64</b>
Marc Dupuis, Ludovic Bonneau, Toshihiko Kawano, T-16	
<b>A Relativistic Symmetry in Nuclei and Effective Field Theory .....</b>	<b>66</b>
Joseph N. Ginocchio, T-16	
<b>Sensitivity of Nuclear Transition Frequencies to Temporal Variations in the Fine Structure Constant (<math>\alpha</math>-dot) .....</b>	<b>68</b>
Anna C. Hayes, James L. Friar, T-16	
<b>Nuclear Masses for Simulation Databases .....</b>	<b>70</b>
Peter Möller, Arnold J. Sierk, T-16	
<b>Prompt Neutrons and Fission Modes .....</b>	<b>72</b>
Patrick Talou, Toshihiko Kawano, T-16	
<b>Novel Heavy Flavor Suppression Mechanisms in the Quark-gluon Plasma .....</b>	<b>74</b>
Ivan Vitev, T-16	
<b>Organic Semiconductor Light-emitting Diodes for Solid State Lighting .....</b>	<b>76</b>
Ping Yang, Enrique R. Batista, Sergei Tretiak, Richard L. Martin, T-12; Avadh Saxena, Darryl L. Smith, T-11	
<b>The Largest Superfluid Gaps in Nature and in the Laboratory .....</b>	<b>78</b>
Sanjay Reddy, Alexandros Gezerlis, Joseph Carlson, T-16	
<b>From Nuclei to Cold Atoms and Back .....</b>	<b>80</b>
Ionel Stetcu, T-16; Bruce R. Barrett, Ubirajara van Kolck, Univ. of Arizona; James P. Vary, Iowa State Univ.	



## Chemistry and Bioscience

<b>Atomic Description of Alzheimer's Amyloid Beta Oligomers .....</b>	<b>84</b>
S. Gnanakaran, T-10; Ruth Nussinov, National Cancer Institute; Angel E. Garcia, Rensselaer Polytechnic Institute	
<b>Automatic Classification of HCV and HIV-1 Sequences with the Branching Index .....</b>	<b>86</b>
Peter Hrabec, Carla Kuiken, Mark Waugh, Shaun Geer, William J. Bruno, Thomas Leitner, T-10	
<b>Recent Vaccination Improves the Early Control of Virus in SHIV Infection but not the Clinical Outcome .....</b>	<b>88</b>
Leslie L. Chavez, Alan S. Perelson, Ruy M. Ribeiro, T-10; Miles P. Davenport, Univ. of New South Wales; John W. Shiver, Lynda G. Tussey, Kara S. Cox, Margaret Bachinsky, Fubao Wang, Lingyi Huang, William A. Schleif, Mary-Ellen Davies, Aimin Tang, Danilo R. Casimiro, Merck Research Laboratories	
<b>Towards a Theory of Transport in Disordered Organic Polymers: Poly(p-Phenylene Vinylene) (PPV) .....</b>	<b>90</b>
Ping Yang, Enrique R. Batista, Sergei Tretiak, Richard L. Martin, T-12; Avadh Saxena, Darryl L. Smith, T-11	



<b>Theoretical Study of Ligand Exchange Reaction on Bis(Imido) Uranium(VI) .....</b>	<b>92</b>
Ping Yang, Enrique R. Batista, T-12; Liam P. Spencer, Brian L. Scott, James M. Boncella, MPA-MC	

<b>Theoretical Studies of Carbon-hydrogen Bond Activation Chemistry of Actinide Complexes.....</b>	<b>94</b>
Ping Yang, Ingolf Warnke, Richard L. Martin, P. Jeffrey Hay, T-12	

<b>Covalency Trends in Group IV Metallocene Dichlorides: Chlorine K-edge X-ray Absorption Spectroscopy and Time Dependent-density Functional Theory.....</b>	<b>96</b>
Ping Yang, Enrique R. Batista, P. Jeffrey Hay, Richard L. Martin, T-12; Stosh A. Kozimor, Carol J. Burns, C-NR; Kevin S. Boland, Laura E. Wolfsberg, C-IIAC; David L. Clark, Daniel E. Schwarz, ADSMS; Steven D. Conradson, Juan S. Lezama, MST-8; Marianne P. Wilkerson, C-ADI; Christin N. Christensen, New Mexico State Univ.	

<b>Spin-orbit Coupling Effects on the Electronic Structure of <math>\text{UO}_2</math>.....</b>	<b>98</b>
Lindsay E. Roy, Richard L. Martin, T-12; Tomasz Durakiewicz, John J. Joyce, MPA-10; Gustavo Scuseria, Rice Univ.	

<b>Spin-orbit and Jahn-Teller Effects in the Excited States of Platinum II Porphyrins.....</b>	<b>100</b>
Cristian V. Diaconu, Richard L. Martin, Enrique R. Batista, T-12; Darryl L. Smith, T-11; Brian K. Crone, Ian H. Campbell, MPA-11; Scott A Crooker, MPA-NHMFL	

<b>Chronic Beryllium Disease: Structural, Spectroscopic, and Reactive Properties of Aqueous Beryllium Ion .....</b>	<b>102</b>
S. Gnanakaran, T-10; Brian L. Scott, MPA-MC	

### Climate Modeling

<b>The LANS-alpha Turbulence Parameterization in Ocean-climate Modeling.....</b>	<b>106</b>
Matthew W. Hecht, Mark R. Petersen, Beth A. Wingate, CCS-2; Darryl D. Holm, CCS-2, Imperial College, London	

<b>A Review of North Atlantic Ocean Modeling in an Eddying Regime .....</b>	<b>108</b>
Matthew W. Hecht, CCS-2	


<b>Advanced Numerical Techniques and Algorithms to Study Abrupt Climate Change .....</b>	<b>110</b>
Balu Nadiga, CCS-2	

<b>Ocean Modeling and the Representation of Unresolved Scales .....</b>	<b>112</b>
Balu Nadiga, CCS-2	

<b>New Result Shows Separation of Time Scales in Rotating and Stratified Flows at High Latitudes .....</b>	<b>114</b>
Beth Wingate, CCS-2; Miranda Holmes, New York Univ; Mark Taylor, Sandia National Laboratories	

<b>Climate, Ocean, and Sea-Ice Modeling.....</b>	<b>116</b>
Philip Jones, Elizabeth Hunke, Bill Lipscomb, Mathew Maltrud, Todd Ringler, Bin Zhao, T-3; Scott Elliott, Matthew W. Hecht, Balu Nadiga, Nicole Jeffery, Mark R. Petersen, Wilbert Weijer, Beth A. Wingate, CCS-2	

<b>Biogeochemistry in Los Alamos Climate, Ocean, and Sea Ice Modeling.....</b>	<b>118</b>
Scott Elliott, CCS-2; Mathew Maltrud, Elizabeth Hunke, T-3	




## Information Science and Technology

<b>RaveGrid: Raster to Vector Graphics for Image Data.....</b>	<b>122</b>
Sriram Swaminarayan, CCS-2; Lakshman Prasad, ISR-2	
<b>Effective and Efficient Visualization and Analysis of Large-scale Data on Emerging Computer Architectures .....</b>	<b>124</b>
Patrick S. McCormick, James Ahrens, CCS-1; Jeff Inman, HPC-1	
<b>Radio-frequency Interference Cancellation for Wideband Communications.....</b>	<b>126</b>
Christopher M. Brislawn, CCS-3; Karl Hill, New Mexico State Univ.	
<b>A Recent Discovery in Anomaly Detection .....</b>	<b>128</b>
Don Hush, Patrick Kelly, Ingo Steinwart, Clint Scovel, CCS-3	
<b>Data Mining in Radiation Portal Monitoring.....</b>	<b>130</b>
Tom Burr, Jim Gattiker, Kary Myers, CCS-6; George Tompkins, D-6	
<b>Framework for Responding to Network Security Events (FRNSE) .....</b>	<b>132</b>
Justin “J.D.” Doak, Aaron Lovato, CTN-5	
<b>Combining MEG and fMRI for Single-pass MEG Analysis – A Simulation Study .....</b>	<b>134</b>
Doug Ranken, HPC-1	
<b>Adaptive and Dependable Server I/O Networking Support in a High Performance Cluster Computing Environment .....</b>	<b>136</b>
Hsing-Bung Chen, HPC-5	
<b>Petascale Data Storage Institute .....</b>	<b>138</b>
Gary Grider, HPC-DO; Carolyn Davenport, James Nunez, HPC-5	



## Materials

<b>Statistical Models for Stockpile and Complex System Health Assessment.....</b>	<b>142</b>
Alyson Wilson, Mike Hamada, Christine Anderson-Cook, Todd Graves, Scott Vander Weil, Aparna Huzurbazar, Benjamin Sims, Greg Wilson, CCS-6	
<b>Models and Simulations of UO<sub>2</sub> Fuel Materials Properties.....</b>	<b>144</b>
Marius Stan, CCS-2; Sven Rudin, John Wills, T-1; Blas P. Uberuaga, Steven M. Valone, MST-8; Shenyang Hu, Pacific Northwest National Laboratory; Petrica Cristea, Univ. of Bucharest	
<b>A 3D Uranium Casting Simulation from Preheat to Cooldown.....</b>	<b>146</b>
Erin I. Barker, Sharen Cummins, CCS-2	
<b>Simulations of Heat and Oxygen Transport in a Nuclear Fuel Element .....</b>	<b>148</b>
Marius Stan, CCS-2; Bogdan Mihaila, MST-6; Alek Zubelewicz, T-3; Juan C. Ramirez, EXPONENT Inc.; Petrica Cristea, Univ. of Bucharest	
<b>The Analysis of Self-diffusion and Migration of Spheres in Nonlinear Shear Flow Using a Traction-corrected Boundary Element Method.....</b>	<b>150</b>
Alan Graham, INST-OFF; Shihai Feng, MST-7; Marc Ingber, Univ. of New Mexico; Lisa Mondy, Sandia National Laboratories	






<b>Simulations of Metallic Nuclear Fuel Casting.....</b>	<b>152</b>
David Korzekwa, Deniece Korzekwa, MST-6; Neil Carlson, Bryan Lally, CCS-2; Mark Christon, Simulia	
<b>A High-accuracy Limited-range Equation of State for Liquid Sodium .....</b>	<b>154</b>
Eric Chisolm, T-1	
<b>Molecular Modeling of Alkaline Fuel Cell Membranes.....</b>	<b>156</b>
Shaji Chempath, Lawrence R. Pratt, T-12; Brian R. Einsla, Clay S. Macomber, James M. Boncella, Jonathan A. Rau, Bryan S. Pivovar, MPA-11	
<b>Accelerated Molecular Dynamics at the Solid-liquid Interface .....</b>	<b>158</b>
Danny Perez, Arthur F. Voter, T-12	
<b>Direct Transformation of Vacancy Voids to Stacking Fault Tetrahedral .....</b>	<b>160</b>
Blas P. Uberuaga, Richard G. Hoagland, Steven M. Valone, MST-8; Arthur F. Voter, T-12	
<b>Heterogeneous Nucleation and Hot Spots at Grain Boundaries in a RDX Bicrystal.....</b>	<b>162</b>
Marc J. Cawkwell, Thomas D. Sewell, T-14; Kyle J. Ramos, Daniel E. Hooks, DE-9	
<b>Homogeneous Dislocation Nucleation and Anomalous Hardening in RDX .....</b>	<b>164</b>
Marc J. Cawkwell, Thomas D. Sewell, T-14; Kyle J. Ramos, Daniel E. Hooks, DE-9	
<b>Numerical Simulation of Projectile/Armor Interaction Using CartaBlanca .....</b>	<b>166</b>
Xia Ma, Duan Z. Zhang, Paul T. Giguere, Qisu Zou, T-3	
<b>Global-local Analysis of Laminated Composite Plates .....</b>	<b>168</b>
Hashem M. Mourad, Todd O. Williams, Francis L. Addessio, T-3	
<b>Thermal Striping.....</b>	<b>170</b>
David J. Torres, Rick M. Rauenzahn, Manjit S. Sahota, T-3	
<b>Perturbed Continuum Approach to Predict the Dynamic Behavior of Metals .....</b>	<b>172</b>
Alek Zubelewicz, T-3; Jimmy Fung, X-2	
<b>Single Rod Fuel-clad Mechanical Interaction .....</b>	<b>174</b>
Manjit S. Sahota, Rick M. Rauenzahn, David J. Torres, Thomas R. Canfield, T-3; Paul J. Maudlin, T-DO; Ralph A. Nelson, ADWP	
<b>Disjoining Pressure for Non-uniform Thin Films .....</b>	<b>176</b>
Bing Dai, Alan Graham, INST-OFF; Antonio Redondo, T-DO; Leslie G. Leal, Univ. of California, Santa Barbara	
<b>The Mechanism of Surfactant Effects on Drop Coalescence .....</b>	<b>178</b>
Bing Dai, Alan Graham, INST-OFF; Antonio Redondo, T-DO; Leslie G. Leal, Univ. of California, Santa Barbara	



## Modeling Complex Networks

<b>Structure Analysis of Large-scale Networks .....</b>	<b>182</b>
Hristo Djidjev, CCS-3	
<b>Using a Massive Agent-based Model to Study the Spread of Pandemic Influenza .....</b>	<b>184</b>
Phillip D. Stroud, Sara Y. Del Valle, D-3; Susan M. Mniszewski, CCS-3; Jane M. Riese, Stephen J. Sydoriak, HPC-1	
<b>Growth, Sustainability, Security, and the Pace of Life in Cities.....</b>	<b>186</b>
Luis M. A. Bettencourt, T-7	
<b>Predicting Mass Evacuation in the Wake of Natural or Human-caused Disasters .....</b>	<b>188</b>
Leticia Cuéllar, Anders A. Hansson, Nick Hengartner, CCS-3; Deborah Kubicek, D-6	
<b>Evaluation of First-born Program for Los Alamos National Laboratory (LANL) Foundation .....</b>	<b>190</b>
Aparna V. Huzurbazar, CCS-6	







## Funding Acknowledgment Cross Reference

CTN Classified and Unclassified Computer Security Program .....	132
Con Alma Foundation .....	190
Department of Defense.....	126, 142, 162, 164, 168
Department of Energy.....	66, 80, 86, 142, 150, 162, 164, 168, 176, 178
Department of Energy, Global Nuclear Energy Partnership.....	144, 148, 152, 154, 166, 170
Department of Energy, National Nuclear Security Administration.....	2, 10, 14, 22, 24, 136, 160, 162, 164
Department of Energy, National Nuclear Security Administration, Advanced Simulation and Computing Program .....	2, 4, 6, 8, 10, 12, 28, 30, 32, 70, 72, 122, 124, 146
Department of Energy, Office of Biological and Environmental Research, Climate Change Prediction Program.....	106, 108, 110, 112, 114, 116, 118
Department of Energy, Office of Nuclear Energy.....	144, 148, 152, 174
Department of Energy, Office of Science.....	62, 64, 68, 74, 114, 172
Department of Energy, Office of Science, High Energy Physics Program .....	56
Department of Energy, Office of Science, Innovative and Novel Computational Impact on Theory and Experiment Program .....	48
Department of Energy, Office of Science, Nuclear Physics Program.....	64
Department of Energy, Office of Science, Office of Advanced Scientific Computing Research.....	34, 36, 38, 40, 42, 44, 46, 124, 150, 158, 176, 178
Department of Energy, Office of Science, Office of Basic Energy Sciences .....	18, 76, 90, 92, 94, 96, 98, 100, 156, 160
Department of Energy, Office of Science, Office of Fusion Energy Sciences .....	36
Department of Energy, Office of Science, Scientific Discovery through Advanced Computing Program.....	48, 58, 64, 78, 116, 118, 138
Department of Homeland Security.....	130, 182, 184, 188
European Union, Information Society as a Complex System Project.....	186
Frederick Reines Postdoctoral Fellowship .....	96
Glenn T. Seaborg Institute.....	92, 94, 96
James S. McDonnell Foundation .....	88

J. Robert Oppenheimer Fellowship.....	74
Los Alamos National Laboratory Foundation .....	190
Los Alamos National Laboratory Center for Nonlinear Studies.....	106
Los Alamos National Laboratory Directed Research and Development Program .....	2, 10, 16, 18, 22, 24, 26, 30, 32, 46, 68, 70, 74, 78, 84, 102, 106, 110, 112, 128, 134, 158, 162, 164, 186
Los Alamos National Laboratory Science and Technology Base Program .....	162, 164
MIND Research Network.....	134
National Aeronautics and Space Administration.....	30, 52, 54
National Institutes of Health.....	86, 88, 134
National Science Foundation .....	80, 186
New Mexico State University and Los Alamos National Laboratory Memorandum of Understanding.....	126
Office of Naval Research .....	162, 164
Thaw Charitable Trust.....	186
United States Nuclear Data Program .....	64

## Author Cross Reference

### A

Adessio, Francis L. ....	168
Ahrens, James ....	124
Akhadov, Elshan A. ....	16
Albright, Brian J. ....	10
Alder, Berni J. ....	30, 32
Anderson-Cook, Christine ....	142

### B

Bachinsky, Margaret.....	88
Baker Jr., George A. ....	62
Barber, John L. ....	30, 32
Barker, Kevin J. ....	6
Barker, Erin I. ....	146
Barrett, Brian ....	12
Barrett, Bruce R. ....	80
Batista, Enrique.....	76, 90, 92, 96, 100
Bergen, Benjamin.....	10
Bettencourt, Luis M. A. ....	186
Boland, Kevin S. ....	96
Boncella, James M. ....	92, 156
Bonneau, Ludovic.....	64
Bowers, Kevin J. ....	10
Brislawn, Christopher M. ....	126
Bruno, William J. ....	86
Burns, Carol J. ....	96
Burr, Tom.....	130

### C

Campbell, Ian H. ....	100
Canfield, Thomas R. ....	174
Carlès, Pierre ....	30
Carlson, Joseph.....	78
Carlson, Neil.....	152
Casimiro, Danilo R. ....	88
Castain, Ralph.....	12
Cawkwell, Marc J. ....	162, 164
Chacón, Luis ....	34, 36
Chavez, Leslie L. ....	88
Chempath, Shaji.....	156

Chen, Hsing-Bung.....	136
Chisolm, Eric.....	154
Christensen, Christin N. ....	96
Christon, Mark ....	152
Chung, Yeo-Jin.....	36
Clark, David L.....	96
Conradson, Steven D. ....	96
Cox, Kara S. ....	88
Cristea, Petricia.....	144, 148
Crone, Brian K. ....	100
Crooker, Scott A. ....	100
Cuéllar, Leticia ....	188
Cummins, Sharen.....	146

### D

Dai, Bing ....	176, 178
Daniel, David.....	12
Daughton, William ....	10
Davenport, Miles P. ....	88
Davenport, Carolyn.....	138
Davies, Mary-Ellen.....	88
Davis, M. Kei ....	6
Del Valle, Sara Y. ....	184
Delzanno, Gian L. ....	36
Desai, Nehal.....	2
Diaconu, Cristian V. ....	100
Djidjev, Hristo.....	182
Doak, Justin “J.D.” ....	132
Dupuis, Marc ....	64
Durakiewicz, Tomasz.....	98

### E

Einsla, Brian R. ....	156
Elliott, Scott.....	116, 118

### F

Feng, Shihai.....	150
Finn, John M. ....	36
Friar, James L. ....	68
Fryer, Chris L. ....	52
Fung, Jimmy.....	172

**G**

Garcia, Angel E.....	84
Gattiker, Jim.....	130
Geer, Shaun .....	86
Germann, Timothy C.....	8, 30, 32
Gezerlis, Alexandros .....	78
Giguere, Paul T.....	166
Ginocchio, Joseph N.....	66
Gnanakaran, S.....	84, 102
Graham, Alan.....	150, 176, 178
Graves, Todd.....	142
Grider, Gary.....	138
Gyrya, Vitally.....	38

**H**

Habib, Salman.....	56
Hamada, Mike .....	142
Hansson, Anders A.....	16, 188
Hay, P. Jeffrey .....	94, 96
Hayes, Anna C.....	68
Hecht, Matthew W. ....	106, 108, 116
Heger, Alexander .....	58
Heitmann, Katrin .....	56
Hengartner, Nick.....	188
Henning, Paul J.....	14
Hill, Karl.....	126
Hoagland, Richard G. ....	160
Hoisie, Adolfy.....	6
Holm, Darryl D.....	106
Holmes, Miranda .....	114
Hooks, Daniel E.....	162, 164
Hraber, Peter.....	86
Hu, Shenyang.....	144
Huang, Lingyi .....	88
Huang, Zhibin .....	30
Hunke, Elizabeth.....	116, 118
Hush, Don.....	128
Huzurbazar, Aparna V. ....	142, 190

**I**

Ingber, Marc.....	150
Inman, Jeff .....	124

**J**

Jeffery, Nicole.....	116
Jones, Philip .....	116
Joyce, John J.....	98

**K**

Kadau, Kai .....	8, 30, 32
Kawano, Toshihiko.....	64, 72
Kelley, Timothy M.....	14
Kelly, Patrick .....	128
Kerbyson, Darren J. ....	6
Korzekwa, David.....	152
Korzekwa, Deniece .....	152
Kozimor, Stosh A.....	96
Kubicek, Deborah.....	188
Kuiken, Carla.....	86
Kurien, Susan.....	44, 46
Kwan, Thomas J.....	10

**L**

Lally, Bryan.....	152
Lang, Mike .....	6
Lapenta, Giovanni.....	36
Leal, Leslie G.....	176, 178
Leitner, Thomas .....	86
Lezama, Juan S.....	96
Lichtner, Peter.....	48
Lipnikov, Konstantin.....	38, 40, 42
Lipscomb, Bill .....	116
Livescu, Daniel.....	2, 22, 24, 26
Lovato, Aaron .....	132
Lu, Chuan .....	48
Lunasin, Evelyn.....	44

**M**

Ma, Xia.....	166
Macomber, Clay S.....	156
Maltrud, Mathew .....	116, 118
Martin, Richard L. ....	76, 90, 94, 96, 98, 100
Maudlin, Paul J.....	174
McCormick, Patrick S.....	124

Mihaila, Bogdan .....	148
Mniszewski, Susan M. ....	184
Mohd-Yusof, Jamaludin .....	2, 22, 24
Möller, Peter .....	70
Mondy, Lisa .....	150
Moulton, David .....	40, 48
Mourad, Hashem M. ....	168
Myers, Kary .....	130

## N

Nadiga, Balu .....	110, 112, 116
Nelson, Ralph A. ....	174
Nunez, James .....	138
Nussinov, Ruth .....	84

## P

Pakin, Scott .....	6
Perelson, Alan S. ....	88
Perez, Danny .....	18, 158
Petersen, Mark R. ....	2, 22, 24, 106, 116
Philip, Bobby .....	34, 48
Pivovar, Bryan S. ....	156
Prasad, Lakshman .....	122
Pratt, Lawrence R. ....	156
Primeau, Francois .....	110

## Q

Quinlan, Dan .....	4
--------------------	---

## R

Ramirez, Juan C. ....	148
Ramos, Kyle J. ....	162, 164
Ranken, Doug .....	134
Rasmussen, Craig E. ....	4
Rau, Jonathan A. ....	156
Rauenzahn, Rick M. ....	170, 174
Reddy, Sanjay .....	78
Redondo, Antonio .....	176, 178
Reed, Darren .....	56
Ribeiro, Ruy M. ....	88
Rickett, Christopher .....	4
Riese, Jane M. ....	184

Ringler, Todd .....	116
Rockefeller, Gabriel .....	52
Rosenblatt, Charles .....	30
Roy, Lindsay E. ....	98
Rudin, Sven .....	144

## S

Sahota, Manjit S. ....	170, 174
Sancho, Jose C. ....	6
Saxena, Avadh .....	76, 90
Schleif, William A. ....	88
Schwarz, Daniel E. ....	96
Scott, Brian L. ....	92, 102
Scovel, Clint .....	128
Scuseria, Gustavo .....	98
Sewell, Thomas D. ....	162, 164
Shashkov, Mikhail .....	42
Shipman, Galen .....	12
Shiver, John W. ....	88
Sierk, Arnold J. ....	70
Sims, Benjamin .....	142
Smith, Leslie .....	46
Smith, Darryl L. ....	76, 90, 100
Sottile, Matthew .....	4
Spencer, Liam P. ....	92
Stan, Marius .....	144, 148
Steinwart, Ingo .....	128
Stetcu, Ionel .....	80
Stroud, Phillip D. ....	184
Svyatskiy, Daniil .....	40, 42
Swaminarayan, Sriram .....	8, 18, 122
Sydoriak, Stephen J. ....	184

## T

Talou, Patrick .....	72
Tang, Aimin .....	88
Taylor, Mark .....	110, 114
Teuscher, Christof .....	16
Tompkins, George .....	130
Torres, David J. ....	170, 174
Tretiak, Sergei .....	76, 90
Tussey, Lynda G. ....	88

**U**

Uberuaga, Blas P. .... 144, 160

**V**

Valone, Steven M. .... 144, 160  
 Van Kolck, Ubirajara ..... 80  
 Vander Wiel, Scott ..... 142  
 Vary, James P. .... 80  
 Vassilevski, Yuri ..... 42  
 Vitev, Ivan ..... 74  
 Voter, Arthur F. .... 18, 158, 160

**W**

Wallstrom, Timothy C. .... 28  
 Wang, Hsing-Lin ..... 16  
 Wang, Fubao ..... 88  
 Warnke, Ingolf ..... 94  
 Warren, Michael S. .... 54, 56  
 Waugh, Mark ..... 86  
 Weijer, Wilbert ..... 116  
 Weiss, M. .... 58  
 White, Martin ..... 56  
 Wilkerson, Marianne P. .... 96  
 Williams, Todd O. .... 168  
 Wills, John ..... 144  
 Wilson, Alyson ..... 142  
 Wilson, Greg ..... 142  
 Wingate, Beth A. .... 46, 106, 114, 116  
 Wofford, John ..... 54  
 Wolfsberg, Laura E. .... 96  
 Woosley, Stan ..... 58

**Y**

Yang, Ping ..... 76, 90, 92, 94, 96  
 Yin, Lin ..... 10  
 Young, Patrick A. .... 52  
 Yu, Huidan ..... 26

**Z**

Zhang, Duan Z. .... 166  
 Zhao, Bin ..... 116  
 Zou, Qisu ..... 166  
 Zubelewicz, Alek ..... 148, 172

# Los Alamos National Laboratory Organizational Abbreviation Legend (used in this publication only)

**Associate Directorate for Stockpile Manufacturing and Support**  
ADSMS

**Associate Directorate for Weapons Physics**  
ADWP

## **Chemistry Division**

C-ADI Advanced Chemical Diagnostics and Intrumentation  
C-IIAC Inorganic, Isotope and Actinide Chemistry  
C-NR Nuclear and Radiochemistry

## **Computer, Computational, and Statistical Sciences Division**

CCS-1 Computer Science for High-Performance Computing  
CCS-2 Computational Physics and Methods  
CCS-3 Information Sciences  
CCS-6 Statistical Sciences  
CCS-DO Division Office

## **Computing, Telecommunications, and Networking Division**

CTN-DC Departmental Computing (CTN-1, -2, and -3)  
    CTN-1 Departmental Computing  
    CTN-2 Departmental Computing  
    CTN-3 Departmental Computing  
CTN-4 Telecommunications  
CTN-5 Network Engineering  
CTN-DO Division Office

## **Decision Application Division**

D-6 Risk Analysis and Decision Support

## **Dyanmic and Energetic Materials Division**

DE-9 Shock and Detonation Physics

## **High-Performance Computing Division**

HPC-1 Scientific Software Engineering  
HPC-2 Computing Operations and Support  
HPC-3 High-Performance Computing Systems  
HPC-5 High-Performance Systems Integration  
HPC-DO Division Office

## **International Space and Response Division**

ISR-1 Space Science and Applications  
ISR-2 Space and Remote Sensing  
ISR-4 Space Instrumentation Systems

## **Los Alamos National Laboratory Institutes Office**

INST-OFF Engineering Institute

## **Materials Science and Technology Division**

MST-6 Metallurgy  
MST-7 Polymers and Coatings  
MST-8 Structure/Property Relations

## **Materials Physics and Applications Division**


MPA-10 Condensed Matter and Thermal Physics  
MPA-11 Sensors and Electrochemical Devices  
MPA-CINT Center for Integrated Nanotechnologies  
MPA-MC Materials Chemistry  
MPA-NHMFL National High Magnetic Field Laboratory

## **Theoretical Division**

T-1 Equation of State and Mechanics of Materials  
T-3 Fluid Dynamics  
T-4 Atomic and Optical Theory  
T-6 Theoretical Astrophysics  
T-7 Mathematical Modeling and Analysis  
T-8 Elementary Particles and Field Theory  
T-10 Theoretical Biology and Biophysics  
T-11 Condensed Matter and Statistical Physics  
T-12 Theoretical Chemistry and Molecular Physics  
T-13 Complex Systems  
T-14 Explosives and Organic Materials  
T-15 Plasma Theory  
T-16 Nuclear Physics  
CNLS Center for Nonlinear Studies  
T-DO Division Office

## **Applied Physics Division**

X-1 -PTA Applied Science and Methods Development-Plasma Theory  
    and Applications  
X-2 Design I





# Advanced Computational Architectures

Computational platforms are the tools of modern science. Computer science itself is centered on creating new methods and technologies and making computers accessible to scientists to enhance their creativity and pace of discovery. The ongoing development of advanced computational architectures is born out of a need for computational speed along with the ever-present driver of “Moore’s Law.” This “law” was actually a statement made in 1965 by Gordon Moore, cofounder of Intel, that the number of transistors on a microprocessor would double approximately every 18 months. Packing more and more technology into smaller and smaller packages translates into rapidly increasingly powerful computational architectures and, theoretically at least, application performance. Four decades later, we are seeing fundamental limits to taking practical advantage of Moore’s Law as measured by realizing actual dramatic performance improvements to applications that run on the new computational engines. Advanced computational architectures seek new, sometimes exotic, architectural designs that accelerate the performance of applications into new realms, and require increased application or software complexity to take advantage of the new designs. Such advances have already reached everyone’s desktops in the form of dual- and quad-core processors, which will soon become homogeneous and heterogeneous many-core computational engines, outpacing software and tools available on such systems.

The revolutionary Roadrunner platform, built with PlayStation 3 chips at its core, is an instantiation of an advanced heterogeneous architecture using standard CPUs and “accelerator” chips simultaneously. Unleashing the power of such platforms for science requires

fundamental thought about existing and new algorithms, numerical methods, and next-generation software tools. The papers in this section speak to the development of new algorithmic techniques and the extension of existing applications onto the Roadrunner platform.

The researchers write about their work developing new language and compiler techniques to make heterogeneous and homogeneous many-core architectures of the future accessible to scientific endeavors. They present results of simulations of new physical phenomena on the Roadrunner, and follow-on, platforms. In addition, they describe tools, methods, and expertise that allow the characterization of existing and theoretical platforms, which have been developed uniquely at LANL. Their enthusiasm for the design of new platforms and applications to truly enhance the performance and throughput of computing issues of interest to the Nation is evident.

Finally, future and emerging computing paradigms are discussed that go well beyond silicon-based platforms into the world of cognitive computing, biocomputation, and nanotechnology. These exciting new and powerful tools will dramatically advance our insight into virtually all issues of scientific and societal importance for the future. LANL is poised to not just take advantage of these new technological breakthroughs, but to create them.

# Direct Numerical Simulation of Turbulent Flows on Roadrunner

Jamaludin Mohd-Yusof, Daniel Livescu, Mark R. Petersen, CCS-2;  
Nehal Desai, CCS-1

**T**urbulence is an inherently multiscale phenomenon, encompassing a large range of length and time scales, all of which must be resolved in a direct numerical simulation (DNS). The technique requires that the numerical errors remain small and not affect the solution. DNS allows a degree of control of specific physical phenomena not accessible in experiments, leading to improved understanding and, ultimately, to models for large multiphysics codes. However, the numerical conditions significantly restrict the range of flows and parameters accessible, as they require very large grids. The current progress in computer hardware, based on heterogeneous architectures, could significantly advance our state-of-the-art turbulence simulations and, thus, turbulence theory and modeling. The IBM Cell Broadband Engine (Cell BE) used in the Roadrunner computer is an example of the newest generation of accelerators. Common to many of these architectures is the existence of multiple processing cores (synergistic processing units, or SPUs, in this case) with independent, high-bandwidth access to local memory.

CFDNS is a Fortran-structured grid code that simulates the compressible Navier-Stokes equations in 3D. Multiple species, each with realistic material properties equation of state (EOS) are allowed, as well as Cartesian, cylindrical, and spherical grid geometries. The bulk of the time in the code is spent in either computing the update equations or computing the derivatives of the various fields. The update equation is computed in a pointwise manner throughout the domain and is straightforward to vectorize once all the terms have been evaluated. The derivative step uses a compact finite difference (Pade) scheme that requires a tridiagonal inversion along the direction in which the derivative is being performed. Since the data is laid out in 3D, this requires a strided load of the data (Fig. 1).

A subset of the base code was ported to the Cell BE architecture and implemented in C with SPU extensions. At present this version only allows periodic boundary conditions, with ideal gas EOS. The derivative calculation vectorizes naturally in the x- and y- directions, since the data is contiguous in z. For the z-derivative, the

data must be reordered in local store to achieve good performance, but the data is not reordered in main memory. Each SPU works on a different portion of the data and executes its own direct memory access (DMA) calls independently, to maximize performance. Similar data partitioning is used for the update step.

The Cell BE implementation of the code was benchmarked against a functionally identical C code running on the Opteron. Significant speedups of 40x ( $64^3$ ) to 50x ( $128^3$ ) were observed. The primary benefit for the Cell BE implementation is the ability to independently access main memory, allowing an effective memory access rate of  $\sim 12$  GB/s (50 percent of theoretical peak) [1]. Further optimization is in progress.

A full parallel implementation of the code is in progress, in anticipation of the full Roadrunner deployment, which will allow simulations at resolutions of  $4096^3$ . These will be by far the largest compressible turbulence simulations ever performed, and we expect them to allow significant advancement on the theory and modeling of compressible turbulence.

**For more information contact Jamaludin Mohd-Yusof at [jamal@lanl.gov](mailto:jamal@lanl.gov).**

[1] J. Mohd-Yusof, N. Desai, D. Livescu, "Roadrunner: Compressible Turbulence Simulation," presented at SC07 (International Conference for High Performance Computing, Network, Storage and Analysis) (2007).

## Funding Acknowledgments

- Department of Energy, National Nuclear Security Administration
- Department of Energy, National Nuclear Security Administration, Advanced Simulation and Computing Program
- Los Alamos National Laboratory Directed Research and Development Program




Fig. 1. Depiction of the division of work between SPUs for the x- and y- derivative calculation. Since each derivative requires data along the entire line in the direction of the derivative being taken, different partitioning is required for each direction. The z-direction derivative and the update equation can use either partitioning.

# Fortran Development Tools: Providing a Roadmap for Application Development on Advanced Computer Architectures

Craig E. Rasmussen, Christopher Rickett, CCS-1; Dan Quinlan, Lawrence Livermore National Laboratory; Matthew Sottile, Univ. of Oregon

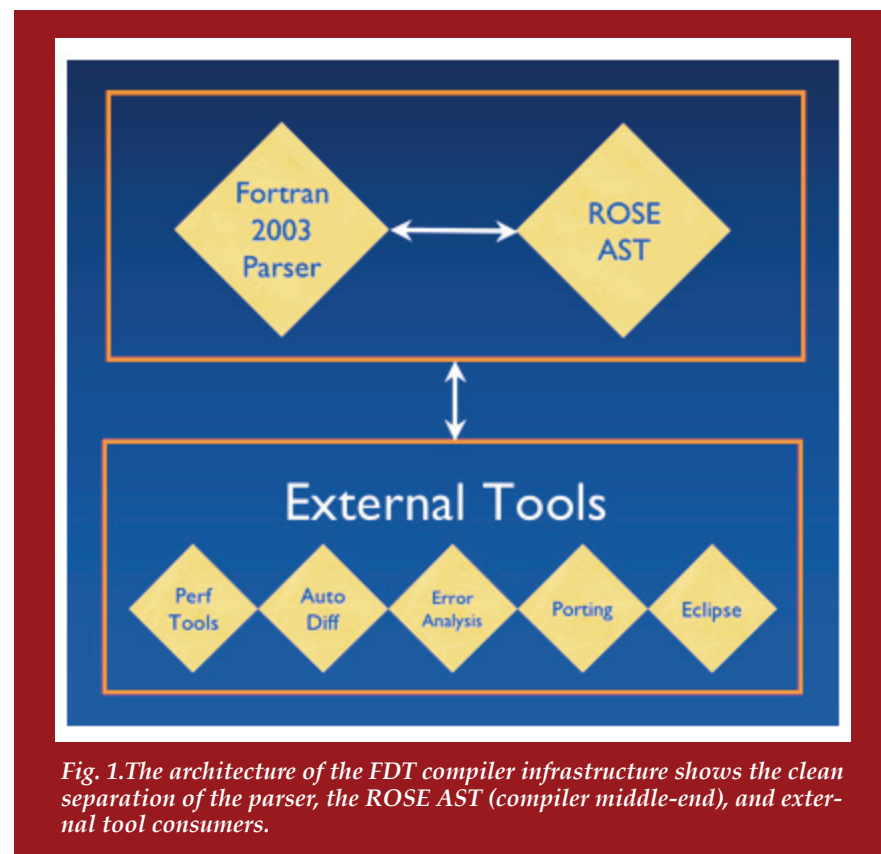
The Fortran Development Tools (FDT) project is dedicated to delivering productivity-enhancing and program-correctness tools to the scientific application developer. The FDT project also provides an important vehicle for source-to-source transformations, enabling research in the mapping of high-level language constructs onto advanced computer architectures like Roadrunner. The goal of this research is to allow the programmer to “write once,” while relying on the source-to-source compiler to target the disparate variety of computer architectures that are available.

Because of the increasing complexity of scientific applications and the computing hardware on which they run, there is a need for tools that support the entire life cycle of application development. These tools include those that support the production and maintenance of code, tools that aid in improving the quality of code, and tools that help developers to both understand and improve application performance. This document describes work begun on a Fortran open-source compiler infrastructure created to support tool development; an infrastructure designed from the beginning for program analysis rather than back-end code generation.

In addition, we also describe a solution to a special need facing today’s application developers in the DOE complex. The need is for transformational tools that support the migration of important DOE applications to run efficiently across a wide range of new computer architectures, including those with multiple homogeneous (and especially heterogeneous) processing elements. Tools are needed to reduce the costs associated with transforming software, and perhaps more importantly, to guarantee that any changes made won’t affect the output of the software.

**Compiler Infrastructure.** The FDT project has created an open Fortran 2003 parser (see <http://sourceforge.net/projects/fortran-parser/>). This parser offers five advantages: 1) it is based on a language grammar (ANTLR) that allows

project outsiders to relatively easily understand and modify the parser to add new language constructs, 2) it is open source, 3) it has a clean separation of the parser from middle-end consumers and Java and C interfaces through which anyone can implement custom abstract syntax trees (ASTs) and tools, 4) it is fully Fortran 2003 compliant, and 5) it has been integrated with the Lawrence Livermore National Laboratory (LLNL) ROSE, C, and C++ infrastructure (see Fig. 1).



*Fig. 1. The architecture of the FDT compiler infrastructure shows the clean separation of the parser, the ROSE AST (compiler middle-end), and external tool consumers.*

**Program Transformations.** Fortran is an important high-level language because it is the language of scientific computation and also because it has many data-parallel constructs. Data parallelism is a high-level abstraction that is, at the same time, both easier to program and gives the compiler more leeway in retargeting a program to different computer architectures.

For example,

$$A = B + s * C \quad (1)$$

is a data-parallel assignment statement where A, B, and C are arrays, and s is a scalar. Note that no explicit iteration over array indices is needed and that the individual operators, plus, times, and assignment are applied by the compiler to each individual array element independently. Thus the compiler is able to spread the computation in (1) across any hardware thread under its control.

While (1) is a very simple example, complete and very concise and elegant programs can be built up with similar statements and intrinsic functions like the array constructors (CSHIFT, MERGE, TRANSPOSE, ...), the array location functions (MAXLOC and MINLOC), and the array reduction functions (ANY, COUNT, MINVAL, SUM, PRODUCT, ...). In addition, with MPI implementations of these intrinsic functions, programs written in this data-parallel subset of Fortran are implicitly parallel; all inter-processor communication is accomplished by these library routines, not by the application programmer.

Microsoft has demonstrated (via the Accelerator project) that scientific applications written in a data-parallel language can run on traditional processors and achieve up to a 17 times speedup on graphical processing unit (GPU) processors, without recoding. The precise amount of speedup obtained depends on the scientific application and the algorithms used.

We are evaluating the use of the FDT/ROSE compiler infrastructure to transform the LANL Pagosa application to run on Roadrunner and other advanced architectures. The Pagosa code was chosen for the evaluation because it uses the data-parallel constructs of Fortran discussed above. The FDT source-to-source compiler was modified to target the IBM Cell Broadband Engine (Cell BE) processor and tested on a small portion of the Pagosa code. Input to the compiler was Pagosa Fortran, and output was C language code employing calls to the IBM ALF library and Cell BE vector intrinsics.

The results of the initial evaluation were promising. The compiler was able to automatically parallelize and vectorize the chosen section of Pagosa code. Run time was ten times faster on the Cell BE processor than on a single-core Opteron. While these initial results were promising, a larger section of Pagosa (employing the MPI CSHIFT library routines) must be tested before the evaluation is complete.

**FDT Project Goals.** The goals of the FDT project in the immediate future are to complete the coupling of our Fortran 2003 parser with the LLNL ROSE compiler

infrastructure and then to develop tools using this infrastructure. In addition, we intend to continue our research in program transformations for advanced computer architectures. This research will provide guidance to developers on ways to program for portability across disparate computer architectures.

**For more information contact Craig E. Rasmussen at [crasmussen@lanl.gov](mailto:crasmussen@lanl.gov).**

### Funding Acknowledgments

- Department of Energy, National Nuclear Security Administration, Advanced Simulation and Computing Program

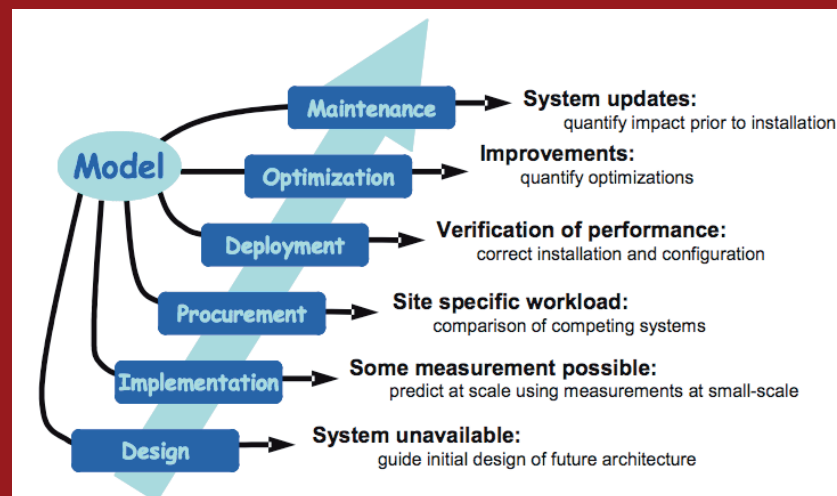
## Performance of Roadrunner under a Realistic Application Workload

Kevin J. Barker, M. Kei Davis, Adolfo Hoisie, Darren J. Kerbyson, Mike Lang, Scott Pakin, Jose C. Sancho, CCS-1

**P**erformance modeling is an important capability developed and refined by the Performance and Architecture Lab (PAL) over the last few years that has been applied to numerous applications and systems at Los Alamos, within the tri-Lab community, and beyond. Using performance models, systems that are not currently installed and hence not available for measurement can be analyzed and their achievable performance can be quantified. Similarly, the performance of applications running on these systems can be analyzed, predicted with high accuracy, and optimized.

Performance modeling has been used throughout the life-cycle of many systems from first design through to maintenance, as shown in Fig. 1. Models have been used to steer early system designs, for instance in collaboration with IBM in the Defense Advanced Research Projects Agency (DARPA) High Productivity Computer Systems (HPCS) program on their Power 7 system (to be available in 2010). Models were used during the procurement of Roadrunner to compare the performance from different system proposals. Models have also been used to analyze systems during installation and, as evident from the ASCI Q installation, they can be used to identify performance issues and assist in system optimizations.

A PAL performance model analyzes both application and system characteristics. Application characteristics are defined uniquely for each application and include processor flow, data structures used, frequency of use and mapping onto the system, and their potential for resource contention. System characteristics include node configuration (processors per node, shared resources), and interprocessor communication (latency, bandwidth, topology). Many of these are measured (for an existing system) or need to be specified/simulated (for a future system). A separate performance model is developed for each application of interest and thus the approach is application-centric.




*Fig. 1. Performance models have been used throughout the life-cycle of many systems.*

The performance of Roadrunner has been analyzed using PAL's highly accurate modeling techniques. PAL developed models of four applications to analyze the performance of Roadrunner prior to its configuration being finalized, and prior to its installation at Los Alamos. This aided in the finalization of the system configuration, and in providing expected performance at large processor-counts for assessment activities. Optimized versions of several Los Alamos applications were used in this process including: Sweep3D (deterministic transport), VPIC (particle-in-cell), SPaSM (molecular dynamics), and Milagro (nondeterministic transport). Sweep3D for the IBM Cell Broadband Engine (Cell BE) was developed by PAL, while the others by researchers in CCS-2.

Many different performance studies were undertaken in analyzing the performance of Roadrunner. For instance, a comparison between the use of Roadrunner in hybrid mode (using both the Cell BEs and Opterons) vs using only the Opterons is shown in Fig. 2 when using only a single node, a single CU (180 nodes), and all CUs. The performance improvement in hybrid mode is over six times for VPIC, over five times for Milagro, approximately four times for Sweep3D on full system and over nine times on a single-node, and over two times for SPaSM. It should be noted in this

performance comparison that the optimization versions of the codes are still being refined and hence the performance improvements may be higher by the time of system installation at Los Alamos.



*Fig. 2. Performance comparison of using the Roadrunner vs using only AMD Opterons.*

Further performance comparisons have shown that the performance of Roadrunner when using the Cell BE will be significantly higher than the performance seen on ASCI Q. A further performance study has shown that the Roadrunner should also achieve higher performance, between 1 and 2.4 times higher, than a system using state-of-the-art AMD or Intel quad-core processors (four sockets per node).

**For more information contact Adolfo Hoisie at [hoisie@lanl.gov](mailto:hoisie@lanl.gov).**

### References

PAL Roadrunner Report #5, "A Note on the Application Performance of the eDP version of the Cell," Los Alamos National Laboratory report LA-CP-07-1147.

K.J. Barker, et al., "Experiences in Scaling Scientific Applications on Current Generation Quad-core Processors," Los Alamos National Laboratory report LA-UR-07-6855.

### Funding Acknowledgments

- Department of Energy, National Nuclear Security Administration, Advanced Simulation and Computing Program



# Exploiting the Cell BE Architecture on Roadrunner for Large-scale Molecular Dynamics

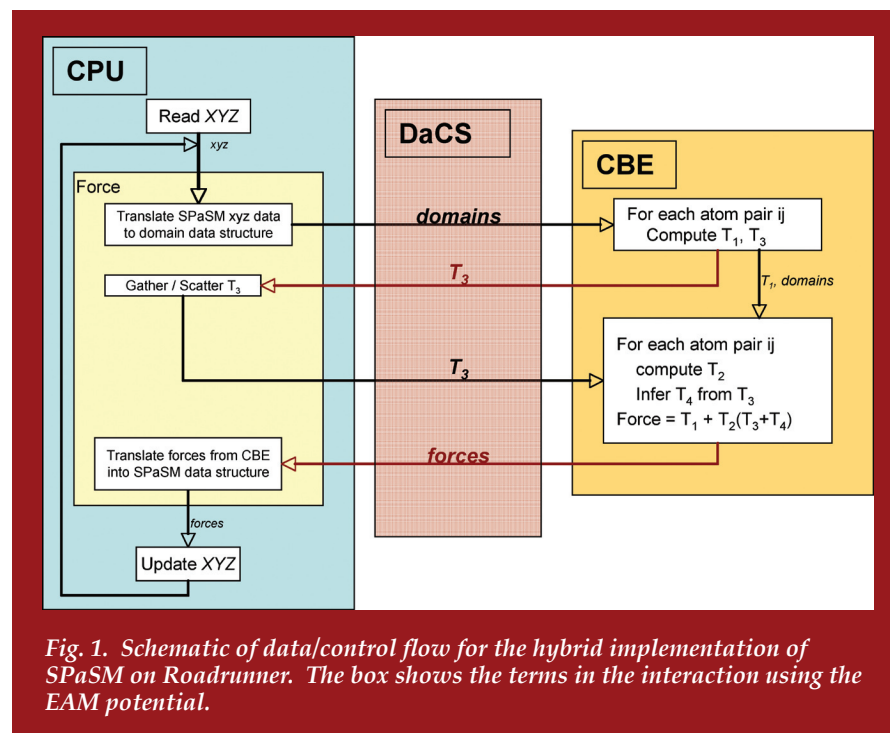
Sriram Swaminarayan, CCS-2; Timothy C. Germann, T-12; Kai Kadau, T-14

Roadrunner, LANL's latest supercomputer due to be delivered at the end of June 2008, is an IBM Cell Broadband Engine (Cell BE) accelerated Opteron cluster consisting of about 7,000 nodes. Each compute node of the machine has two dual-core Opterons and two dual Cell BE blades. The machine is capable of delivering a peak of 1.4 petaflops double precision, almost all of which (>96%) is delivered by the Cell BEs. Consequently, to achieve good performance, any code that is written *must* utilize the Cell BEs. Complicating this simple axiom are the following:

- The Cell BEs are true local accelerators in that they are connected to the host via PCI Express slots and have no network access. Consequently all network communication (e.g., MPI calls) must be routed through the Opterons on the node.
- The endianness of the Cell BE is different from that of the Opteron, and so any data that is transferred between the Opteron and the Cell BE has to be byte-swapped before the other host can interpret it.
- The Cell BE is a vector machine. To exploit its impressive performance one has to vectorize the kernels of code one plans to run on it.
- Although the Cell BE has 8 GB of main memory, the vector units (called synergistic processing units, or SPUs) have only 256 kB each, called the Local Store, and the user has to explicitly manage the data residing in this local memory.

In this paper we briefly describe our strategy for overcoming each of the above points to port SPaSM, a large-scale massively parallel molecular dynamics (MD) code developed at LANL, to the Roadrunner machine.

The first step in accelerating SPaSM was to identify the computational hot spots. In SPaSM, as in most MD codes, over 95 percent of the time is spent in the force computation, making the selection of the hot spot easy. To compute forces, SPaSM decomposes space into subdomains and computes interactions between atoms in neighboring subdomains. Thus, once we receive boundary subdomains from



neighboring processors, the computation is entirely local. Our model therefore was to have the Opterons manipulate the positions of the particles while the Cell BEs computed the forces. This eliminates the need for the Cell BEs to do any communication with any other Cell BE. This model is schematically represented in Fig. 1, which shows the flow of control and data during the course of a simulation. The data that is sent to the Cell BE is endian, converted on the Opteron before transmission, and the forces returned are endian, converted on the Cell BE before transmission. We found the time spent in endian conversion to be an insignificant fraction of the time spent computing forces.

On the Cell BE itself, the data for the atoms (about 200 MB for a million atoms) is resident in main memory. To compute forces, the atoms are *streamed* through the Cell BE one domain at a time. This allowed us to stay within the 256 kB Local Store while still being able to compute forces on all atoms. For our code the Local Store



on the SPU is approximately utilized as follows: 30 kB for the force code, 70 kB for the embedded atom method (EAM) potential tables (three tables of 3000 doubles each), and the rest available for use as a cache for atomic positions and forces.

As a first cut we have implemented a hybrid parallel version of this model using full neighbor lists (i.e., we do not exploit Newton's third law to reduce the computation by a factor of two). In spite of this, for system sizes that are larger than 8,000 atoms per Opteron core (we've tested sizes up to a few million atoms per Opteron core), we get a speedup of about three times over the Opteron version. This is particularly encouraging given that not only are we doing twice the work of the Opteron version, we are achieving this speedup on cells that are about 1/8th the speed of that which will be provided with Roadrunner.

In conclusion, we have ported an existing large-scale MD code running a nontrivial potential to the Roadrunner architecture and achieved good performance gains. By the time Roadrunner is delivered, we expect to improve the performance further, not only because the Cell BEs on Roadrunner are eight times faster, but also because we are redesigning the streaming algorithm to use half-neighbor lists. Thus, although at first blush the Roadrunner architecture appears to be difficult to program, we have demonstrated that it is relatively simple to accelerate existing codes and achieve high levels of performance.

**For more information contact Sriram Swaminarayan at [sriram@lanl.gov](mailto:sriram@lanl.gov).**

### **Funding Acknowledgments**

- Department of Energy, National Nuclear Security Administration, Advanced Simulation and Computing Program

## Petascale Plasma Physics Simulations on Roadrunner Using VPIC


Benjamin Bergen, CCS-2; Brian J. Albright, Kevin J. Bowers, William Daughton, Thomas J. T. Kwan, Lin Yin, X-1-PTA

For many problems of interest in kinetic plasma physics, including laser plasma interactions and magnetic reconnection, the particle-in-cell (PIC) method is the only viable simulation tool. Members of the VPIC team are using this method to pioneer discoveries in areas such as inertial confinement fusion (ICF) research at the National Ignition Facility (NIF).

VPIC [1], a highly optimized, state-of-the-art PIC implementation for fully relativistic, 3D, explicit, kinetic plasma simulations, has been specifically designed to minimize data motion, an increasingly important computational consideration, as modern processing architectures consume more time in moving data than in performing computations. Additionally, VPIC aggressively makes use of the short-vector, single instruction multiple data (SIMD) execution pipelines available on many of these architectures.

Recent work by X-1-PTA and CCS-2 staff has resulted in the adaptation of VPIC to run efficiently on the IBM Cell Broadband Engine (Cell BE) chip, which forms the basis of the new petaflop/s Roadrunner supercomputer that LANL will acquire in 2008. The VPIC particle-push performance on this hardware is staggering, at 0.17 billion cold particle advances per second per Cell BE. This opens the possibility of simulating plasmas on the full Roadrunner system as large as one trillion particles, one billion computational cells, and one million time integrations in 10 days of wall-clock time, a simulation scale 10 times larger than what is feasible today.


VPIC has been applied with success in both basic and applied science. In Fig. 1, we show electron density isosurfaces resulting from magnetic reconnection of a Harris geometry (a plasma equilibrium with a reversal of magnetic field across a neutral layer, with pressure equilibrium established by an initial region of high density at the sheet center) in the case of an intermediate guide field. The simulation follows the dynamics of 16 billion particles on a  $1000 \times 100 \times 1000$  mesh (physical domain  $200 \times 20 \times 200 d_i$ , where  $d_i$  is the ion inertial length). The color indicates magnetic field strength  $B_y$  ( $y$  is the direction of the initial current) at a time when the



*Fig. 1. VPIC simulations of magnetic reconnection in a pair plasma with an intermediate guide field. Displayed are isosurfaces of electron density colored by the magnetic field  $B_y$  ( $y$  is in the initial current and guide field direction). The isosurfaces undulate from interaction of the tearing modes (wave vector along  $x$ ) with the drift-kink modes (along  $y$ ).*

dominant diffusion region has formed and is influenced by a current-aligned instability, the drift-kink mode [2]. This calculation required 36 hours of wall-clock time on 500 AMD Opteron cores.

In Fig. 2, we show a 3D simulation of stimulated Raman scattering (SRS) in the kinetic regime ( $kl_d = 0.34$ , where  $k$  is the initial Langmuir wavenumber and  $l_d$  is the plasma Debye length). The simulation used 18 billion particles on a mesh of size  $2436 \times 486 \times 486$  (a physical domain  $90 \times 18 \times 18$  microns) run on 1008 AMD Opteron cores. The electron plasma waves self-localize to regions of narrow transverse extent, as shown by the filament structures in the longitudinal electric field. SRS saturates via the trapped electron self-focusing of Langmuir waves [3]. SRS is one of the pernicious instabilities that jeopardize the success of realizing inertial confinement fusion ignition at the multi-billion-dollar NIF at Lawrence Livermore National Laboratory.



*Fig. 2. VPIC simulations of stimulated Raman scattering (SRS) in the kinetic regime. Isosurfaces of longitudinal electrostatic field show filament structures resulting from trapped electron self-focusing of Langmuir waves during SRS saturation. The laser is launched from the simulation geometry near face.*

The combined abilities of VPIC to provide excellent science with exceptional computational efficiency make it a natural candidate for adaptation to the Roadrunner hybrid supercomputing architecture. With the development of a working hybrid code, the VPIC team has already made outstanding progress towards ensuring that petascale plasma physics simulations on this exciting new resource become a reality.

For more information contact Brian Albright at [balbright@lanl.gov](mailto:balbright@lanl.gov).

- [1] K. J. Bowers, et al., "Ultra high performance 3D electromagnetic relativistic kinetic plasma simulation," *Phys. Plasmas* (in press).
- [2] L. Yin et al., "Fully kinetic 3D simulations of collisionless reconnection in large-scale plasmas," *Phys. Rev. Lett.* (submitted).
- [3] L. Yin et al., *Phys. Rev. Lett.* **99**, 265004 (2007).

#### Funding Acknowledgements

- Department of Energy, National Nuclear Security Administration
- Department of Energy, National Nuclear Security Administration, Advanced Simulation and Computing Program
- Los Alamos National Laboratory Directed Research and Development Program

# Open MPI and the Roadrunner Project


David Daniel, CCS-1; Ralph Castain, ISR-4; Brian Barrett, Sandia National Laboratories; Galen Shipman, Oak Ridge National Laboratory

Large-scale parallel scientific applications are, with few exceptions, based on the Message Passing Interface (MPI) [1], which is a standardized communication library allowing the coordination of, and data exchange between, the constituent processes of an application.

The Open MPI Project [2] is an open source MPI implementation that is developed and maintained by a consortium of academic, research, and industry partners. Open MPI is therefore able to combine the expertise, technologies, and resources from all across the high-performance computing community with the goal of creating the best MPI library available. Open MPI offers advantages for system and software vendors, application developers, and computer science researchers.

**Open MPI and Los Alamos.** A team in CCS-1 was a founding member of the Open MPI Project and led much of the development through the first several releases, leveraging experience gained in the development and deployment of LA-MPI [3] (the MPI implementation currently in use on the Pink, Flash, and Lightning Linux production systems). Consequently, more than other MPI implementations, Open MPI has been driven by the needs of the LANL computational science community, specifically in the areas of portability, consistent user interface, and performance on systems of interest to the Laboratory.

Of particular importance for the Roadrunner project is the issue of scalability, that is, the ability to run a parallel science application efficiently on thousands of processors. As an example, consider how the memory used by an MPI library scales with the number of processors. Figure 1 compares memory usage of Open MPI 1.0 with another MPI implementation (MVAPICH v0.9.5) on systems using an InfiniBand network [4]. Many MPI implementations are aimed at relatively small clusters—up to a few hundred processors—so concern about the memory used per process is a secondary consideration. The result is that when scaled to thousands of processors an ever-increasing fraction of memory is consumed by the MPI library rather than being available to an application. Not so for Open MPI, which



*Fig. 1. Scaling of memory usage with number of processes in Open MPI and MVAPICH.*

was designed from the beginning to scale well on very large systems, and has been successfully used on Sandia National Laboratories' Thunderbird Linux cluster at scales in excess of 8,000 processes.

Open MPI is currently available on several InfiniBand Linux clusters at LANL including the institutional Coyote cluster, as well as the Roadrunner project phase 1 clusters Yellowrail and Redtail.

**Open MPI and Roadrunner.** As described elsewhere in this paper, the final Roadrunner system is an advanced Linux cluster of thousands of heterogeneous computational nodes, each comprising several AMD Opteron processors teamed with IBM Cell Broadband Engine accelerators. The network too, is heterogeneous, consisting of InfiniBand between computational nodes and a custom PCI Express-based link between Opteron and Cell BE "blades."

Although this is a highly complex and novel architecture, remarkably, Open MPI is sufficiently flexible and portable that prototype systems were quickly made available to application developers investigating algorithmic techniques and performance for Roadrunner. This proved crucial to the assessment phase of the Roadrunner project leading to the decision to go ahead with procurement of the final system.

Open MPI was also used as the prototype platform for the IBM software team developing new programming environments (DaCS and ALF) aimed at simplifying the task of application development for Roadrunner and future heterogeneous computer systems.

Work is currently in progress in collaboration with IBM to improve the scalability and performance of Open MPI for Roadrunner-scale systems. Also planned is improved support for the Roadrunner system architecture so that users can, if they choose, write a complete heterogeneous application using MPI rather than the new DaCS/ALF frameworks.

**For more information contact David Daniel at [ddd@lanl.gov](mailto:ddd@lanl.gov).**

[1] Message Passing Interface Forum, <http://www.mpi-forum.org/>.

[2] Open MPI, <http://www.open-mpi.org/>.

[3] The Los Alamos Message Passing Interface, <http://public.lanl.gov/lampi/>.

[4] G. Shipman et al., electronic *Proceedings 20th IEEE International Parallel and Distributed Processing Symposium*, (2006).

### **Funding Acknowledgments**

- Department of Energy, National Nuclear Security Administration, Advanced Simulation and Computing Program

## Adaptation of Milagro Implicit Monte Carlo to the Roadrunner Hybrid Architecture


Timothy M. Kelley, Paul J. Henning, CCS-2

**I**mplicit Monte Carlo (IMC) is a method of simulating thermal X-ray transport for nonlinear, time-dependent problems such as supernova explosions and inertial confinement fusion. Milagro is a set of radiation-only applications used to drive the Jayenne Project's IMC libraries for method validation and testing.

The Roadrunner supercomputer architecture increases a standard Opteron cluster floating point throughput by an order of magnitude with Cell Broadband Engine (Cell BE) processors. To study the feasibility of running IMC on Roadrunner, we adapted a subset of the Milagro applications and the underlying transport libraries to the Roadrunner hybrid architecture. We measured a substantial speedup of 5x on a representative test problem using prototype hardware.


In order to learn as much as possible about how to use the hybrid Roadrunner design, our study took two approaches. One started with the existing code base and evolved incrementally, first adapting the core particle transport routines to the Cell BE processor, then restructuring the full application (running on the Opterons) to offload particle transport to the Cell BE. A more revolutionary approach freely changed data structures and algorithms, packing data into work blocks that can be processed on any available computer resource—Cell BE or Opteron. Both approaches used interfaces that abstract the accelerator (Cell BE) and decouple the generation of particles (or work blocks) from processing them. These design considerations should enable the changes made for Roadrunner to apply to other emerging architectures as well as conventional clusters.

In Milagro's 3D wedge representation of 2D RZ geometry, we devised a problem consisting of a radiation source propagating through a complex pattern of thick and thin materials (Fig. 1). The thick materials were ratio-zoned to better resolve transport boundary layers. Although the two approaches were pursued independently, both achieved similar speedups on double bend; Fig. 2 shows



*Fig. 1. The mesh used in the double bend test problem.*





*Fig. 2. Preliminary speedup of hybrid IMC on the double bend test problem as a function of time step, measured on prototype hardware. Very early in the simulation, there is not enough work to amortize the cost of running hybrid, but as the material heats up and more particles are used, hybrid quickly becomes more efficient than the single Opteron. By time step 2000, hybrid is five times faster (9.5 s wall clock hybrid, 48.0 s Opteron).*

the speedup per time step. On prototype hardware—without Roadrunner’s enhanced double precision Cell BE processors—we measured a 5x speedup over  $10^4$  time steps (1  $\mu$ s simulation time). While the performance profile may vary for different problems, these results are excellent, and we are working on additional improvements.

In the future, we will merge the best of both approaches into the Jayenne project software base and research revolutionary transport methods enabled by this new architecture.

**For more information contact Tim Kelley at [tkelley@lanl.gov](mailto:tkelley@lanl.gov).**

#### **Funding Acknowledgments**

- Department of Energy, National Nuclear Security Administration

## Future and Emerging Computing Paradigms and Machines


Christof Teuscher, Anders A. Hansson, CCS-3; Elshan A. Akhadov, MPA-CINT; Hsing-Lin Wang, C-PCS

After more than 50 years of exponential computing power growth under Moore's law, sequential computing as we know it is increasingly hitting a wall, mainly because of fundamental physical limits, economical considerations, and reliability issues because of miniaturization. In order to understand, model, predict, and simulate tomorrow's natural and man-made complex systems, we need computer science and computer engineering to keep going at least at the same pace, which requires massive investments in radically new machines, paradigms, and models of computation.

The 21st century promises to be the century of cognitive, bio, nano, and information science [1]. We are interested in future and emerging computing paradigms and machines that will enable us to 1) go beyond the horizon of silicon-based top-down engineered electronics, and 2) open new and previously unexplored application areas and environments for information-processing devices, such as invisible and undetectable electronics based on bio- and nano-molecules or randomly assembled systems that are extremely hard to reverse engineer.

Molecular self-assembly is often mentioned as the holy grail of nanotechnology, which allows us to easily and cheaply build systems of enormous complexity. The grand challenge consists of using this functional and structural complexity for performing efficient communication and computations.

Our work, which combines theory, simulations, and experiments, goes far beyond other approaches [2,3] for novel interconnects. We have taken an extremely fabrication-friendly stance of building novel computing machines by building functional structures that we can very easily and cheaply self-assemble, without having to worry about the device's reliability, regularity, or homogeneity. The first steps solely focused on interconnectivity issues because the interconnect is today's limiting factor of performance on electronic chips.



*Fig. 1. Self-assembled nanowires fabricated at LANL. Image courtesy: Hsing-Lin Wang.*

Nanowires can be grown in various ways using diverse materials such as metals and semiconductors. We have demonstrated that Ag nanowires can be fabricated on top of conducting polymer membranes via a spontaneous electrodeless deposition (self-assembly) method [4,5] (Fig. 1 shows a prototype of such an assembly). We will use such wires to densely interconnect computing nodes. On the theoretical side, we have shown that such irregular, nonclassical, and self-assembled interconnects can have superior performance and robustness over regular assemblies [6,7].

Our interdisciplinary and integrated approach in both the experimental and theoretical aspects opens novel perspectives on future information processing devices. With a better understanding of the relevant design trade-offs, the



fabrication issues, and the communication and computing aspects, our research represents a significant step towards building alternative, more complex, bottom-up fabricated information processing devices in a cheaper and ultimately simpler way. We argue that mastering these future and emerging technologies is critical for the Laboratory to remain a world leader in computing.

**For more information contact Christof Teuscher at [christof@lanl.gov](mailto:christof@lanl.gov).**

- [1] C. M. Roco and S. W. Bainbridge, eds., *Converging Technologies for Improving Human Performance: Nanotechnology, Biotechnology, Information Technology and Cognitive Science* (2002).
- [2] Y. Chen et al., *Nanotechnology*, **14**, 462-468 (2003).
- [3] A. DeHon, *IEEE Transactions on Nanotechnology*, **2**(1), 23-32 (2003).
- [4] C. Teuscher, *Chaos*, **17**(2), 026106 (2007).
- [5] C. Teuscher, *Second International Conference on Nano-Networks, Nano-Net 2007*, Catania, Italy (2007).
- [6] H.-L. Wang, W. Li, E. Akhador, and Q. Jia, *Chemistry of Materials*, **19**, 520 (2007).
- [7] W. Li, Q. Jia, and H.-L. Wang, *Polymer*, **47**, 23 (2006).

### **Funding Acknowledgments**

- Los Alamos National Laboratory Directed Research and Development Program

# Exploiting the Cell BE Architecture on Roadrunner for Long-time Atomistic Dynamics

Danny Perez, Arthur F. Voter, T-12; Sriram Swaminarayan, CCS-2

The last few years have seen the demise of single-core superscalar microprocessor architectures in favor of multicore processors. With physical and manufacturing constraints and power density requirements inhibiting further increases in clock frequency, the industry is now turning to heterogeneous architectures such as multiple different cores and/or specialized vector co-processing units to deliver increased performance. One such next-generation architecture is IBM's Cell Broadband Engine (Cell BE) [1]. A Cell BE is composed of a general-purpose Power PC element (the PPE) to which are connected eight synergistic processing elements (SPEs), which are vector processing units. Taking full advantage of the SPEs can theoretically deliver an impressive 204 GFlop/s in single precision or 102 GFlop/s in double precision. The prospect of such stunning performance has generated a lot of excitement around this new architecture, culminating in the design of the upcoming LANL supercomputer, Roadrunner.

This paradigm shift in processor design poses serious challenges to scientific programmers. Indeed, these potential gains can only be realized at the price of extensive redesigns of existing applications. High-performance applications will now need to manage fine-grained, intra-processor parallelization over a heterogeneous collection of processing units on top of the usual coarser-grained inter-processor parallelization, the former now becoming increasingly performance-critical. New strategies are thus needed if the promises of next-generation supercomputers like Roadrunner are to be achieved.

In this communication, we present one such strategy aimed at achieving high performance in atomistic molecular-dynamics simulations on Cell BE-based supercomputers. We find that a time-based parallelization on systems of modest size using the parallel replica method can outperform the conventional space-based parallelization of large systems through a more efficient use of the SPEs. We demonstrate this using benchmarks on a single Cell BE processor and deduce the expected performance on Roadrunner.

The molecular-dynamics simulation method (MD) is one of the cornerstones on which our understanding of the atomic-scale behavior of materials is built. Traditionally, good performance of large-scale simulations has been tied to the minimization of interprocessor communication. This is typically accomplished by assigning distinct regions of space to different processors. Communication is then required to inform each processor of the current state of the atoms in neighboring regions of space. Maximal performance is achieved when the volume-to-surface ratio of each region is high, keeping the computation-to-communication ratio high. This strategy is efficient for simulating very large systems over relatively short time scales, making it ideal to study problems such as shock wave propagation [2].

An alternative strategy is to exploit a time-based decomposition of the problem using the parallel-replica dynamics method [3]. In this method, developed here at LANL, a complete replica of the system resides on each processor. For dynamics dominated by infrequent transitions between different configurations, one can show that the time required for a single replica to observe a transition is statistically equivalent to the sum of times taken by a number of independent replicas before any one of them observes a transition. One can thus obtain a nearly linear speed-up when increasing the number of processors. Since the replicas are independent, communication is only necessary to transmit the new state of the system each time a transition occurs. As such, communication is minimized for small systems where transitions are infrequent. This method is thus ideally suited to study slow processes, such as defect diffusion and annealing, over very long time scales [4].

Since these two simulation approaches are aimed at different regimes (size scale vs time scale), a given problem typically will be well suited to one and not the other. However, viewed purely from the perspective of efficiency, the parallel replica strategy offers distinct advantages on the Cell BE. The overall efficiency of simulations on Roadrunner-type supercomputers strongly depends on the efficient use of the SPEs. The challenge is formidable since the local memory on the SPEs [the so-called Local Store (LS)] is limited to 256 kB for both instructions and data. Thus, at most a few thousand atoms can be stored in the LS, the rest remaining in the Cell BE's main memory (which is 8 GB on Roadrunner). This implies that if the system assigned to each Cell BE is larger than this, atoms will need to be streamed through the LS of the SPEs a few thousand at a time to keep them continuously fed. Further, neighbor list information can not be cached in the LS and has to be recalculated during each force computation. The resulting overhead reduces the overall efficiency of the computation. The crucial point is that the constraints placed on the space-decomposition scheme by maximizing the coarse-grained volume-to-surface ratio (i.e., minimizing the communication between Cell BEs) reduces the efficiency of the computation of the SPEs since systems will not fit in LS. However, this situation should be ideal for the parallel replica method.




Fig. 1. Time per force computation per atom as one Cell BE processor for two storage strategies.

To test this hypothesis, we implemented force computations (the most time-consuming task in atomistic simulations) corresponding to the two aforementioned strategies on the Cell BE architecture using the embedded atom method (EAM) potential of Voter[5]. The two implementations essentially differ by the storage scheme of the atomic coordinates: in the first (corresponding to large-scale space-decomposition MD), they are stored in the Cell BE's main memory and streamed through the SPE's LS, while in the second (corresponding to the parallel replica method) the entire system is resident in each SPE's LS. The respective timings are reported in Fig. 1, where the time required to compute the force acting on a single atom using a single Cell BE processor is presented as a function of the system size.

The results clearly demonstrate the advantage of storing the system in the SPE's LS. Indeed, the time per force call per atom is around 0.2  $\mu$ s in this case, compared with about 4  $\mu$ s for storage in main memory, an increase in speed by a factor of 20. While precise values can be slightly affected by other parameters of the calculation, our timings clearly show that large performance gains can be achieved by dedicating each Cell BE to a system of small size. As mentioned above, this gain will be offset by communication costs if used in conjunction with a space decomposition strategy but is a perfect fit to the parallel replica method.

Using these results we can extrapolate the performance of a parallel replica code on Roadrunner. As an example, we recently used the parallel-replica method to

study vacancy void dynamics in fcc metals, and discovered a surprising propensity for direct collapse to a stacking fault tetrahedron [4] (see *Direct Transformation of Vacancy Voids to Stacking Fault Tetrahedral*, pg. 160 in this book). If a similar study of post-collision annealing of radiation-produced defects was to be carried out on Roadrunner using 2000 atoms, it should be possible to simulate the evolution of the system at a rate of up to 1 ms/day of computation. Further, the parallel replica method can be combined with other accelerated dynamics methods, most notably the hyperdynamics method [6], to push further out (by anywhere from a factor of 10 to a factor of  $10^5$ ) the timescales that can be reached using atomistic simulations. Thus, Roadrunner will enable the parallel replica method to access experimentally relevant timescales for slow processes like aging or defect annealing studies, while retaining a fully atomistic description of the system.


In conclusion, we have shown that considerable performance gains can be obtained for atomistic simulations on the Cell BE architecture if each Cell BE is dedicated to a system of a few thousand atoms, a regime ideally suited for the parallel replica method. Further, this strategy could also be used to design cheap, commodity-based computers (using Sony's PlayStation 3 game consoles, for example) geared toward long-time molecular-dynamics simulations. Our approach opens the door to efficient long-time simulations on next-generation supercomputers, nicely complementing the existing capabilities of large-scale molecular-dynamics and thus considerably extending the range of applicability of atomistic simulations.

**For more information contact Danny Perez at [danny\\_perez@lanl.gov](mailto:danny_perez@lanl.gov).**

- [1] *Cell Broadband Engine Architecture*, IBM Systems and Technology Group, 330 Hopewell Junction, NY 12533-6351, 1st ed. (2007).
- [2] K. Kadau, et.al, *Phys. Rev. Lett.* **98**, 135701 (2007).
- [3] A.F. Voter, *Phys. Rev. B* **57**, R13985 (1998).
- [4] B.P. Uberuaga, et al, *Phys. Rev. Lett.* **99**, 135501 (2007).
- [5] A.F. Voter, Los Alamos National Laboratory Technical Report LA-UR-93-3901, Los Alamos National Laboratory (1993).
- [6] A.F. Voter, *Phys. Rev. Lett.* **78**, 3908 (1997).

#### Funding Acknowledgments

- Department of Energy, Office of Science, Office of Basic Energy Sciences
- Los Alamos National Laboratory Directed Research and Development Program





# Applied Mathematics, Fluid Dynamics, and Magnetohydrodynamics

The articles in this section are examples of the type of research carried out in the Theory, Simulation, and Computation Directorate in the techniques of applied mathematics and hydrodynamics, and the study of hydrodynamics in the presence of magnetic fields. These articles are brief summaries of original research in statistical inference, solvers for magnetohydrodynamics, numerical grid and mesh generation, turbulence, instabilities in fluid systems (including fluid mixing), flows in porous media, diffusion, and reactive and stratified flows.

The work spans the gamut from theoretical formulations to model development to implementation into large-scale simulations. The applications for these techniques are also varied. For example, turbulence is an area of fluid dynamics important in gas engines, combustion, and astrophysics, as well as laboratory experiments. Turbulence in conjunction with shock waves can also lead to instabilities at the interface that separates two fluids. These instabilities are typically difficult to solve by numerical methods but these articles offer some approaches capable of tackling the problem accurately and efficiently.

In the articles on applied mathematics, for example, there is a short summary of work on statistical methods of inference using Bayesian statistics. Researchers in our directorate describe how they have tackled the numerical solution of the equations of extended magnetohydrodynamics, a highly coupled set of nonlinear partial differential equations. Other methods to numerically solve partial differential equations require the generation of grids that can be adapted to the particular geometry of the problem. One such technique, presented in one of the articles in this section, is based on optimization approaches developed for a completely different problem.




## Simulations of Forced Isotropic Compressible Turbulence

Mark R. Petersen, Daniel Livescu, Jamaludin Mohd-Yusof, CCS-2

Compressible turbulent flows arise in many science and engineering applications, including conventional gas engines, supersonic combustion ramjet engines, and astrophysical jets, and are also important for many laboratory applications. Many phenomena in compressible turbulence are currently poorly understood compared with incompressible turbulence: the mechanism for energy transfer, interaction among scales and universality, intermittency, and the effects of compressibility on the turbulent energy spectrum. Of particular interest are shocklets, which are short-lived, localized shocks produced by the fluctuating fields of the turbulent eddies [1]. Shocklets convert turbulent kinetic energy to internal energy via the pressure-dilatation correlation, and increase kinetic energy decay through dilatational dissipation (dilatation is the compressible component of the velocity).

We use direct numerical simulations (DNS) to simulate compressible forced isotropic turbulence and to quantify the effects of compressibility in turbulent flows [2]. DNS resolves all scales down to the viscous dissipation range, but on today's computers is restricted to small, idealized domains. As DNS offers a wealth of information inaccessible in laboratory experiments, statistics from DNS studies are critical to parameterize turbulence models, e.g., unclosed terms in Reynolds-averaged Navier-Stokes (RANS) and moment closures or subgrid stresses in large eddy simulations (LES).

Isotropic turbulence simulations are traditionally forced at the smallest wavenumber using spectral numerical methods. We have implemented a new forcing technique that is efficient for (nonspectral) compact finite difference numerical methods [3]. The successful implementation of this physical-space forcing is critical for the efficient use of large computer platforms like Roadrunner, where communications among thousands of nodes impose important computational challenges. The standard domain is a 3D periodic box (Fig. 1). Production runs are being started with resolutions of  $1024^3$  and higher. The effects of compressibility become more significant with increasing rms Mach number (the ratio of the root-mean-squared fluctuating velocity to the speed of sound (Fig. 2)). Of particular interest is the influence of the equations of state (both pressure and caloric) on the turbulence



*Fig. 1. Pressure in a forced isotropic compressible turbulence simulation. Compressible turbulence is characterized by regions with large pressure gradients where shocklets may develop.*




Fig. 2. Statistics that vary with the Mach number show how increased compressibility affects the simulation. The equilibrium averaged values of kinetic energy (left) and dilatational energy (right) increase with Mach number. Here  $E_d$  is the energy in the dilatational (compressible) velocity component, and  $E_s$  is that in the solenoidal (incompressible) component.

properties, never considered before, but important in laboratory applications. Statistics collected from high-resolution turbulence simulations will quantify the energetic effects of compressibility and shocklets. These simulations also produce the turbulent inlet condition for simulations of turbulence-shock wave interaction [4].

For more information contact Mark Petersen at [mpetersen@lanl.gov](mailto:mpetersen@lanl.gov).

- [1] R. Samtaney, et al., *Phys. Fluids*, **13**, 1415-1430 (2001).
- [2] D. Livescu, *J. Fluid Mech.* **450**, 35-66 (2002).
- [3] C. Rosales and C. Meneveau, *Phys. Fluids*, **17**, 5106 (2005).
- [4] J. Mohd-Yusof, et al., "Direct Numerical Simulation of Turbulent Flows on Roadrunner" (in this publication pg 2).

#### Funding Acknowledgments

- Department of Energy, National Nuclear Security Administration
- Los Alamos National Laboratory Directed Research and Development Program



# Shock-turbulence Interaction

Jamaludin Mohd-Yusof, Daniel Livescu, Mark R. Petersen, CCS-2

**T**he interaction between a shock wave and an interface separating two different fluids can generate a wide variety of fluid motions. The description of the interaction combines compressible phenomena, such as shock interaction and refraction, and hydrodynamic instability, including nonlinear growth and transition to turbulence. In general, a perturbation initially present at the interface is amplified following the refraction of the shock. This class of problems is usually referred to as the Richtmyer-Meshkov (RM) instability. Accurately solving all the scales of the problem imposes severe limitations on both the spatial and temporal resolutions and, to date, no direct numerical simulation (DNS) in which all the scales of the problem are accurately resolved has been performed for the fully compressible case. However, with the current advance of supercomputers it is now possible to perform DNS of the RM instability generated by weak shocks, without the simplifications used in previous simulations. Due to various factors, including sensitivity to initial conditions, wavelength-dependent growth rates and the effects of physical and numerical dissipation (in non-DNS codes), there is considerable difficulty in understanding the differences between experimental and numerical results. Therefore the effects of molecular transport properties on the development of the instability remain outstanding open questions, which will be addressed with the current simulations.

In the first part of the project, we consider the interaction between a turbulent field and multiple passes of a shock wave for a single fluid. While this is a simpler problem than the multifluid shock-turbulence interaction, it still addresses several open questions, such as the variation of turbulence properties in the presence of an anisotropic production mechanism [1], and modeling approaches for shocked turbulence. There have been several previous studies of this simplified problem at low resolution [2,3]. Our study brings some improvements in the numerical approach and will represent the largest simulations to date of such a configuration.


The simulations are performed using the CFDNS compressible turbulence simulation code, with periodic boundary conditions in the two spanwise directions. The inlet BC is supersonic with a turbulence velocity field generated from a tri-periodic simulation. These inlet generation simulations are performed with the

correct background velocity imposed in order to remove the ‘frozen turbulence’ approximation that would otherwise be required and was used in the previous studies [2,3]. That is, the inlet fields are snapshots of an advecting forced turbulence field at different times, rather than spatial slices of a single field at a fixed time, as done in the previous studies.

Special care is taken to ensure that spurious waves due to the outlet boundary do not affect the simulation results. We use an artificial acceleration field to convect the outgoing perturbations supersonically out of the domain, even though the flow there is actually subsonic. This minimizes the errors due to both physical and numerical waves generated at the outflow boundary.

At the time of writing, final testing of the code improvements is underway, prior to beginning the production runs. An example of the results of these tests is shown in Fig. 1.

**For more information contact, Jamaludin Mohd-Yusof at [jamal@lanl.gov](mailto:jamal@lanl.gov) and Daniel Livescu at [livescu@lanl.gov](mailto:livescu@lanl.gov).**



- [1] D. Livescu and C.K. Madnia, *Phys. Fluids*, **1698**, 2876 (2004).
- [2] S. Lee, S.K. Lele, P. Moin, *J. Fluid Mech.* **251**, 562 (1993).
- [3] K. Mahesh, S.K. Lele, and P. Moin, *J. Fluid Mech.* **334**, 379 (1997).

#### Funding Acknowledgments

- Department of Energy, National Nuclear Security Administration
- Los Alamos National Laboratory Directed Research and Development Program

# Convergence Effects on Acceleration-driven Instability

Huidan Yu, Daniel Livescu, CCS-2

**T**he Rayleigh-Taylor instability (RTI) [1] occurs at an interface between two fluids having different densities when the fluids are subject to acceleration (gravity) pointing from the heavy to the light. In many applications, RTI occurs in convergent geometries (e.g., cylindrical and spherical), which can include implosion and explosion. Examples are implosion in inertial confinement fusion (ICF), laser-induced launching of a flyer plate, stellar pulsations, and supernova explosions. Geometrical convergence or divergence introduces effects not present in Cartesian geometry. In literature, fundamental instability studies and high-resolution, high-accuracy computations of these flows in convergent geometries are scarce.

In many situations, perturbation growth in cylindrical or spherical geometry also includes a contribution from the interface movement, which is called the Bell-Plesset effect [2,3]. Both interface movement and acceleration-driven mechanisms concur and contribute to the perturbation growth towards turbulence.

There is a complex phenomenology associated with the evolution of the interfacial instability. If the initial perturbations are extremely small, the early growth of the instability is exponential with time. It can be analyzed using the linearized dynamic equations. When the amplitude of the perturbation grows, spikes and bubbles are formed due to nonlinear interactions among initial perturbations of different frequency. In the nonlinear growth stage, these spikes and bubbles compete, amalgamate, and entrain with each other. The motions become increasingly complex. Eventually, the flow becomes turbulent.

We have studied RTI-driven flow in cylindrical geometry in the linear early stage [4]. In parallel, we have developed an accurate direct numerical simulation (DNS) high-performance computation code to compute all stages of the instability growth, from early linear through weakly nonlinear and nonlinear growth to fully turbulent mixing.

The configurations examined in the linear stage include 3D cylindrical as well as 2D axisymmetric and circular unperturbed interfaces, which are compared with the Cartesian case with planar interface (PI). Focuses are on the effects of compressibility and geometrical convergence (or divergence) on the instability growth and the differences between implosion (gravity acting inward) and explosion (gravity acting outward). Compressibility can be characterized by two parameters: 1) static Mach number based on isothermal sound speed, and 2) ratio of specific heats, which in general have opposite influence, stabilization and destabilization, on the instability growth, similar to the Cartesian case [5]. Instability is found to grow faster in the 3D cylindrical case than in the Cartesian case in implosion but slower in explosion. In general, the difference between implosion and explosion is profound for the cylindrical cases but marginal for the planar interface. The influences of Atwood number, interface radius, and the transition between the 3D and 2D circular interface cases have also been considered.

We use the CFDNS code to perform direct numerical simulations of the compressible Navier-Stokes equations and species transport equations in cylindrical geometry for a Rayleigh-Taylor configuration. In addition to the parameters examined during the early stages, we are also interested in the influence of the equation of state and molecular transport properties (viscosity, heat conduction, and mass diffusion) on the instability development. Large-resolution numerical simulations are underway.

**For more information contact Huidan Yu at [hyu@lanl.gov](mailto:hyu@lanl.gov) and Daniel Livescu at [livescu@lanl.gov](mailto:livescu@lanl.gov).**

- [1] D.H. Sharp, *Physica D*, **12**, 3 (1984).
- [2] G.I. Bell, "Taylor instability on cylinders and spheres in the small amplitude approximation," Los Alamos Scientific Laboratory report LA-1321 (1951).
- [3] M.S. Plesset, *J. Appl. Phys.*, **25**, 96 (1954).
- [4] H. Yu and D. Livescu, "Linear analysis of Rayleigh-Taylor instability between immiscible compressible fluids in cylindrical geometry," submitted to *Phys. Fluids*, (2008).
- [5] D. Livescu, *Phys. Fluids*, **16**, 118 (2004).

### Funding Acknowledgment

- Los Alamos National Laboratory Directed Research and Development Program

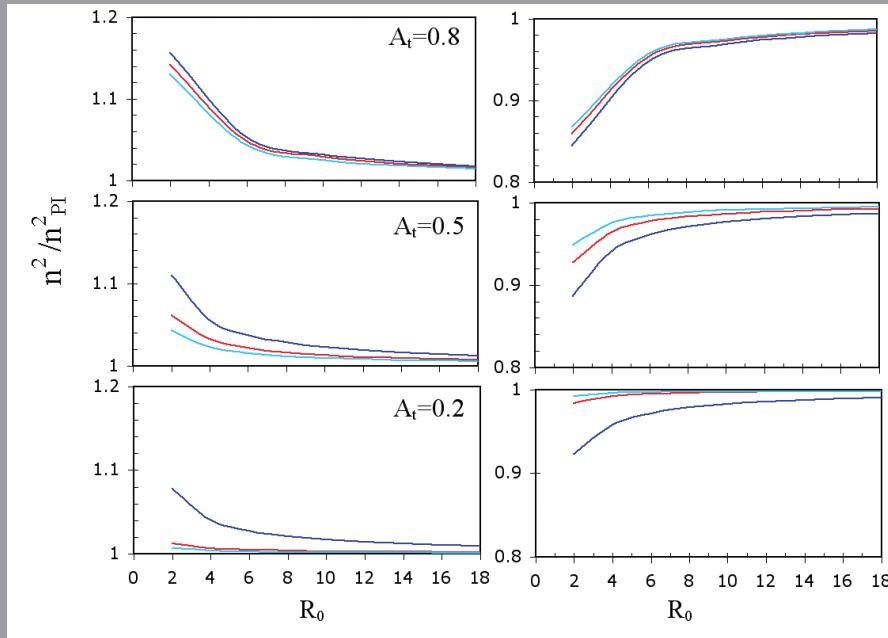


Fig. 1. Dependence of normalized growth rate of  $n$  on normalized position of the interface  $R_0$ . Left column: implosion; right column: explosion. The interface in implosion is less stable than in explosion. The differences between the cylindrical and Cartesian growth rates are largest at higher  $At$  numbers and small  $R_0$ .

# The Marginalization Paradox

Timothy C. Wallstrom, T-13

Statistical methods of inference are increasingly important to many areas of Laboratory interest. Whether the problem is to assess the safety and reliability of the nuclear stockpile, to determine the causes of global warming, or to estimate the phylogenetic tree of the HIV virus, uncertainty comes into play and needs to be treated intelligently.

Bayesian inference has recently experienced a resurgence of interest and is heavily used in many Laboratory applications. Nevertheless, there are some striking paradoxes afflicting Bayesian inference as it is frequently practiced. In recent work, I have clarified a key paradox of Bayesian inference, the Marginalization Paradox (MP) of Dawid, Stone, and Zidek.

In the Bayesian framework, we have a *statistical model*,  $p(x|\theta)$ , which describes the probability of observing data  $x$  when the true parameter is  $\theta$ , and a prior probability  $\pi(\theta)$ , which describes our uncertainty about  $\theta$  before we see the data. Bayes' law allows us to infer the posterior,  $\pi(\theta|x)$ , which describes the new probability of  $\theta$  after seeing the data:

$$\pi(\theta|x) \propto p(x|\theta)\pi(\theta).$$

The MP and some related paradoxes arise from the use of improper prior distributions, which is to say, prior distributions that integrate to infinity. Such priors arise naturally when we wish to discuss ignorance about parameters on infinite domains, such as the real line. They were used by Laplace, and are still in common use today.

The MP arises in problems in which both the data and parameter can be split into parts:  $\theta = (\eta, \zeta)$  and  $x = (y, z)$ . For problems with appropriate symmetry, the marginal densities,  $\pi(\zeta|y, z)$  and  $p(z|\eta, \zeta)$ , are independent of  $y$  and  $\eta$ , respectively; we may write them as  $\tilde{\pi}(\zeta|z)$  and  $\tilde{p}(z|\zeta)$ , respectively. There are then two different ways to compute the marginal posterior  $\pi(\zeta|z)$ . It turns out that when the prior is improper, the results are often incompatible. The paradox is often dramatized as a conflict between two Bayesians,  $B_1$  and  $B_2$ ; see Fig. 1.




Fig. 1.  $B_1$  and  $B_2$  reach incompatible results for the same problem by marginalizing over different quantities.

I have studied this problem by representing the inference  $\pi(\theta|\xi)$  as the limit of posteriors based on proper priors. I have shown that if the limit is taken in an appropriate sense, introduced by Mervyn Stone in the 1960s, then there is no MP.

The resolution of the MP leads to some surprising consequences. First, it turns out that many problems have no limit. This suggests that the idealization represented by the improper prior does not define a statistically meaningful problem. Second, even when the limit does exist, Bayes' law may give the wrong answer. These important consequences are under further investigation.

For more information contact Timothy Wallstrom at [tcw@lanl.gov](mailto:tcw@lanl.gov).

- [1] T. C. Wallstrom, "The marginalization paradox and probability limit," in J.M. Bernardo et al., ed., *Bayesian Statistics* **8**, 669-674, Oxford University Press (2007).
- [2] T. C. Wallstrom, "The marginalization paradox and the formal Bayes' law," in *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, edited by K. Knuth et al., AIP Conference Proceedings **954**, 93-100 (2007). <http://arxiv.org/abs/0708.1350>, and video of associated talk available on Google video (2007).

#### **Funding Acknowledgment**

- Department of Energy, National Nuclear Security Administration, Advanced Simulation and Computing

# The Importance of Fluctuations in Fluid Mixing

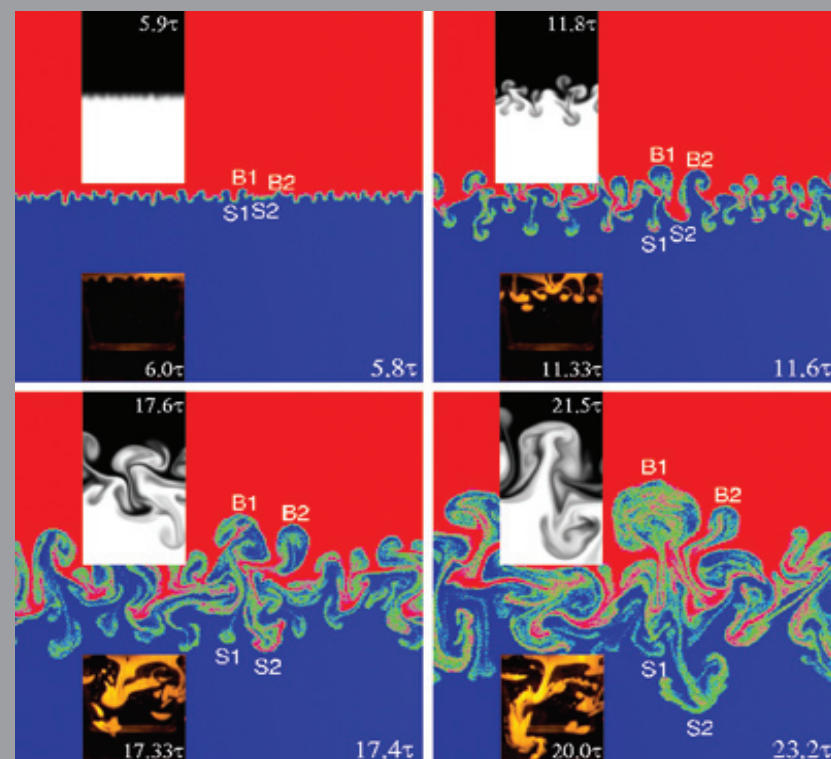
Kai Kadau, John L. Barber, Timothy C. Germann, T-14; Pierre Carlès, Université Paris; Berni J. Alder, Lawrence Livermore National Laboratory; Charles Rosenblatt, Zhibin Huang, Case Western Reserve Univ.

A ubiquitous example of fluid mixing is the Rayleigh-Taylor instability, in which a heavy fluid initially sits atop a light fluid in a gravitational field. The subsequent development of the unstable interface between the two fluids is marked by several stages. At first, each interface mode grows exponentially with time, before transitioning to a nonlinear regime characterized by more complex hydrodynamic mixing. Unfortunately, traditional continuum modeling of this process has generally been in poor agreement with experiment. Here we indicate that the natural, random fluctuations of the flow field present in any fluid, which are neglected in continuum models, can lead to qualitatively and quantitatively better agreement with experiment. We performed billion-particle atomistic simulations with the LANL SPaSM code and magnetic levitation experiments with unprecedented control of initial interface conditions. A comparison between our simulations and experiments reveals good agreement in terms of the growth rate of the mixing front, as well as the new observation of droplet breakup at later times. These results improve our understanding of many fluid processes, including interface phenomena that occur, for example, in supernovae, the detachment of droplets from a faucet, and inkjet printing. Such instabilities are also relevant to inertial confinement fusion (ICF), in which a millimeter-sized capsule is imploded to increase pressure and temperature, and initiate nuclear fusion reactions between deuterium and tritium. An understanding of the flow in such an imploding capsule is necessary to achieve a cleaner source of energy than the ordinary fission reactions used in today's nuclear power plants. Overall, our results suggest that the applicability of continuum models would be greatly enhanced by explicitly including the effects of random fluctuations.

For more information contact Kai Kadau at [kkadau@lanl.gov](mailto:kkadau@lanl.gov).

## Reference

K. Kadau et al., *Proc. Natl. Acad. Sci. USA*, **104**, 7741-7745 (2007).



**Fig.1. Time evolution of RT instability.** Atomistic simulation of a 2D RT instability employing 570 million DSMC particles, where the heavy fluid (red) is on top of the light fluid (blue), and the gravitational field points downwards ( $A=0.67$ ). The initially flat interface roughens due to thermal fluctuations, and modes near the most unstable mode develop first. Later, merger processes become dominant. At even later times, breakups occur that reduce the overall growth of the mixing layer. The position of individual spikes (S1, S2) and bubbles (B1, B2) shows the variation of movement and development of individual spikes and bubbles. Macroscopic magnetic levitation 2D RT experiments (color insets, showing only the spikes in orange, system width 7cm,  $\tau=0.05$  sec,  $A=0.29$ ) exhibit similar breakups. For comparison, a cutout from a 2D continuum simulation is shown (grayscale inset,  $A=0.5$ ) (24). The outer scale Reynolds-numbers at the last frame for the atomistic simulation, experiment, and continuum simulation are 2200, 3000, and 1800, respectively.






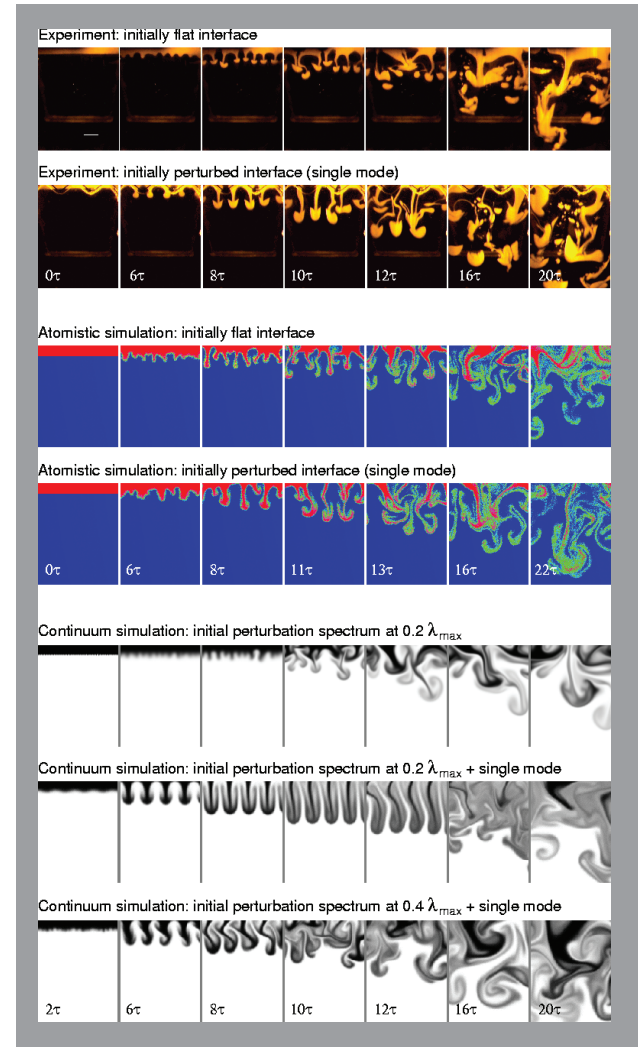
Fig. 2. Spike velocity. Spike velocity as a function of time for atomistic simulation, magnetic levitation experiment, and continuum simulation as shown in Fig.1. The dashed lines indicate the slopes, from which the growth coefficients  $\alpha_S$  are determined. Whereas atomistic simulations and experiments undergo a breakup at  $t = 12\tau$  and  $t = 8\tau$ , respectively, continuum simulations, after an initial phase where the growth is dominated by the initial perturbation spectrum, show no such effect. The horizontal lines indicate the average terminal velocities for atomistic simulations and experiments. The inset shows the instantaneous  $\alpha_S$  value  $\frac{h^2}{4Ag h}$  (23). For all three methods, the penetration depth was measured as the point where the penetrating fluid reached 1 percent or 5 percent number fraction. The experimental data was omitted for  $t > 18\tau$  due to boundary effects. In order to compare the  $\alpha_S$  at different Atwood numbers  $A=0.67, 0.29$ , and  $0.5$  for atomistic simulations, experiments, and continuum simulations, respectively, the corresponding  $\alpha_B$  are calculated (by using the experimental ratio of  $\alpha_S/\alpha_B$ ) to  $0.068$  (atomistic),  $0.055$  (experiments), and  $0.034$  (continuum). Note that  $\alpha_S$  increases with  $A$ , whereas  $\alpha_B$  is constant (16,20).

Fig. 3. Initial conditions. Differences between spikes developing from a flat interface (top row: experiment, horizontal white bar corresponds to 1 cm; third row: atomistic simulation) and from an interface perturbed with a single mode (second row: experiment; fourth row: atomistic simulation).

Experimentally, the interface was given a single-mode perturbation by introducing a pair of sinusoidal, weakly magnetically permeable wires adjacent to the magnetically levitated interface. This perturbation had an amplitude to wavelength ratio  $A_0/\lambda_0=0.0030$ , and the experiment had an Atwood number  $A=0.29$ . The perturbed atomistic simulation ( $A=0.67$ ) had  $A_0/\lambda_0=0.0036$ .

The continuum simulations (24,31) ( $A=0.5$ ) were started with a small multimode perturbation spectrum centered on  $0.2 \lambda_{max}$  (row 5), a small multimode perturbation spectrum centered on  $0.2 \lambda_{max}$  added to a single mode ( $A_0/\lambda_0=0.0036$ , row 6), and a small multimode perturbation spectrum centered on  $0.4 \lambda_{max}$  plus a single mode ( $A_0/\lambda_0=0.0036$ , row 7).

Perturbations initially boost the growth and introduce some order to the flow structure. However, later the influence of perturbations gets weaker. The single-mode perturbation was placed in all cases at  $\sim 2 \lambda_{max}$ . (The atomistic DSMC and continuum simulation frames shown here are a small subset of the whole simulation domain.)



#### Funding Acknowledgments

- Department of Energy, National Nuclear Security Administration, Advanced Simulation and Computing Program
- Los Alamos National Laboratory Directed Research and Development Program
- National Aeronautics and Space Administration

# The Scaling Properties of Atomistic Fluid Dynamics Simulations

John L. Barber, M. Kai Kadau, T-14; Timothy C. Germann, T-12;  
Berni J. Alder, Lawrence Livermore National Laboratory

**F**luid dynamics is usually described by continuum models such as the Navier-Stokes (NS) equation, and such models have had great success in describing real-world phenomena. In addition, the results can often be scaled in various ways to describe an entire class of related systems. However, the validity of continuum models is in practice an emergent phenomenon of the underlying dynamics of the atoms comprising the fluid. In recent years, the available computational capacity has increased to the point where direct simulation of these underlying atomistic dynamics has become possible for increasingly large systems [1]. Nevertheless, there has been considerable resistance to the acceptance of such simulations. In particular, it has been asserted that they do not scale, i.e., that the results of atomistic simulations at the largest scales currently accessible ( $\sim 50 \mu\text{m}$ ) are not relevant to the description of larger, macroscopic systems.

The Navier-Stokes equation with gravity,

$$\rho(\partial_t u + u \cdot \nabla u) = -\nabla p + \mu \Delta u + \rho g$$

is invariant under the transformation  $u \rightarrow u/a$ ,  $x \rightarrow ax$ ,  $t \rightarrow a^2 t$ ,  $g \rightarrow g/a^3$ ,  $p \rightarrow p/a^2$ . In order to determine if atomistic simulations possess a similar scaling, we have performed a series of direct simulation Monte Carlo (DSMC) simulations of the Rayleigh-Taylor instability (RTI), in which a heavy fluid initially lies on top of a light fluid in the presence of gravity. These simulations had various values of the gravitational acceleration  $g$ , and contained up to 5.7 billion particles. If these simulations obey the scaling described above, then expressing the dimensions of the system in terms of a scale length  $\lambda$  and a scale time  $\tau$ , which scale with gravity as  $\lambda \sim g^{-1/3}$  and  $\tau \sim g^{-2/3}$ , will cause the results of different-gravity runs to appear equivalent.

Figure 1 shows a series of snapshots in the development of the RTI for several gravities at various values of  $t/\tau$ . Each image is  $\sim 50 \lambda$  in width. It is clear that these images are qualitatively similar for identical values of  $t/\tau$ . To make this notion more quantitative, we have found that at


$t/\tau = 14$  the fractal dimension  $D$  of the interface is — to within numerical uncertainty — the same for each of these gravities. Unscaled results for the bubble and spike positions as a function of time are shown in Fig. 2. Note that the length and time scales vary by factors of 5 and 26, respectively, between the lowest and highest gravities. Upon scaling by  $\lambda$  and  $\tau$ , however, these results approximately collapse onto one another (Fig. 3). (The moderate disagreement for large  $t/\tau$  can be attributed to the influence of the fluctuations and compressibility present in any atomistic simulation.) Furthermore, a scaled comparison between an atomistic simulation and experimental RTI results (Fig. 3, inset) shows similar agreement, even though the unscaled length and time scales differed by factors of  $42 \times 10^6$  and  $70 \times 10^3$ , respectively, and the gravity differed by a factor of  $10^{10}$ . This suggests that atomistic simulations can, in fact, be scaled up to describe macroscopic situations.

**For more information contact John L. Barber at [jlbarber@lanl.gov](mailto:jlbarber@lanl.gov).**

[1] K. Kadau et al., *Proc. Nat. Acad. Sci.*, **104**, 7741-7745 (2007).

## Funding Acknowledgments

- Department of Energy, National Nuclear Security Administration, Advanced Simulation and Computing Program
- Los Alamos National Laboratory Directed Research and Development Program



**Fig. 1.** Time evolution of the species field for different gravities and Atwood number  $A = 0.67$  (red: heavy particles, blue: light particles, green: mixture). Snapshots at different times and gravities were chosen to show a subset of the system with a width of  $50 \lambda$ . The fractal dimension  $D$  of the interface at  $t/\tau = 14$  is given in the top right corner for each gravity. To within noise, the scaled results exhibit no difference for different gravities.




Fig. 2. Unscaled bubble and spike positions for different values of gravity and  $A = 0.67$ . The ratios between the largest and smallest length and time scale are 5 and 26, respectively.




Fig. 3. Scaled bubble/spike positions for different values of gravity (same data as in Fig. 2). Inset: Scaled simulation and experimental spike positions, for which the unscaled length scales, time scales, and gravities differ by factors of  $42 \times 10^6$ ,  $70 \times 10^3$ , and  $10^{10}$ , respectively.

# Scalable Solvers for 3D Extended Magnetohydrodynamics

Luis Chacón, T-15; Bobby Philip, T-7




Fig. 1. Sample 3D simulations with PIXIE3D, a fully implicit extended MHD simulation code. The left figure depicts densities and velocities in the nonlinear stage of a 3D Kelvin-Helmholtz unstable configuration with differential rotation. The figure on the right depicts temperature during the evolution of a helical perturbation in a 3D toroidal fusion device.

Extended magnetohydrodynamics (XMHD) is a highly coupled, nonlinear partial differential equation system characterized by multiple physical phenomena that span a very large range of length and time scales. Numerically, XMHD results in very ill-conditioned algebraic systems, which require implicit time stepping for efficiency, and spatial adaptivity to resolve dynamically forming thin layers. We summarize our progress towards fully implicit 3D XMHD using Newton-Krylov methods, and towards the implementation of fully implicit adaptive mesh refinement (AMR).

The purpose of this project is to explore efficient time integration algorithms for XMHD. XMHD describes processes occurring in plasmas, which are common in the natural world, and are highly important for many advanced DOE applications. The predictive simulation of plasmas is a crucial capability, one that requires the solution of strongly coupled equations on high resolution, adaptive meshes.

Temporally, XMHD supports multiple, disparate time scales in the form of normal

modes (waves). These waves are generally much faster than the slow dynamical time scales of interest. Explicit techniques are unsuitable for XMHD, as they must resolve the fastest supported time scale. Semi-implicit methods can effectively step over explicit time-step constraints, but concern exists about their accuracy for large time steps [1]. Implicit methods, however, can step over fast time scales to accurately resolve the dynamical time scale [2], but a large-scale system of stiff nonlinear equations must be solved every time step. For this, robust and efficient parallel iterative solvers are necessary. Our work focuses on the development of scalable, fully coupled Newton-Krylov-based (NK) solvers for XMHD.

Preconditioning is key to accelerating NK. Formulating suitable preconditioning strategies for XMHD is a main goal of this project. We aim at enabling the use of multi-level techniques in the preconditioning stage, because they have the potential of delivering optimal convergence rates with problem size [3,4,5]. In this regard, we have pioneered

the “physics-based” preconditioning concept. At the heart of the approach is the “parabolization” of the hyperbolic XMHD system, which recovers block diagonal dominance and renders XMHD multilevel-friendly. Sample results of an XMHD grid convergence study showing the excellent algorithmic performance of the approach are given in Table 1. In particular, note that the CPU is proportional to the number of mesh points, and the number of GMRES iterations is not only bounded, but improves under grid refinement.



Table 1. XMHD grid convergence study

1 time step,  $\Delta t = 1.0$ , V(3,3) cycles,  $mg\_tol=1e-2$ 

Grid	GMRES/ $\Delta t$	$CPU_{exp}/CPU$	$\Delta t/\Delta t_{exp}$
32x32	22	0.44	110
64x64	12	1.4	238
128x128	8	6.2	640
256x256	7	30	3012

We have implemented these developments in a 3D XMHD simulation tool, PIXIE3D [6]. PIXIE3D features a conservative, solenoidal finite-volume discretization in general curvilinear geometry, with remarkable stability properties. Sample PIXIE3D results are depicted in Fig. 1.

Table 2. Implicit AMR grid convergence study.

 $\Delta t = 1$  (fixed),  $\eta_k = 0.1$ ,  $\epsilon_{rel} = \epsilon_{abs} = 10^{-7}$ , 2 SI iterations, V(3,3) cycles

Levels	GMRES its/ $\Delta t$				
	1	2	3	4	5
$32 \times 32$	3.4	7.9	12.0	19.3	33.7
$64 \times 64$	6.5	11.7	19.1	33.2	—
$128 \times 128$	12.5	20.1	27.2	—	—
$256 \times 256$	19.9	27.5	—	—	—
$512 \times 512$	26.3	—	—	—	—

Despite the efficiency benefits of a fully implicit temporal integration of XMHD, implicit time stepping alone is not sufficient for a truly optimal algorithm. Disparate length scales in XMHD result in dynamically evolving thin structures, which need to be resolved as the simulation proceeds. Using uniform grids for this purpose would result in unnecessarily large problems. In order to minimize the number of degrees of freedom required for a given simulation, grid adaptation is required. However, the integration of AMR and fully implicit time stepping is a challenging task. We have addressed this challenge by coupling physics-based preconditioned NK with structured adaptive mesh refinement (SAMR, provided by the SAMRAI package). We have developed an implicit AMR capability on 2D reduced MHD [7], which has demonstrated excellent algorithmic performance (Table 2). A calculation of a tilt instability using implicit AMR is shown in Fig. 2.




Fig. 2. Tilt instability calculation using implicit AMR.

For further information contact Luis Chacón at [chacon@lanl.gov](mailto:chacon@lanl.gov).

- [1] D.A. Knoll et al., *J. Comput. Phys.*, **185**, 583-611 (2003).
- [2] Ibid.
- [3] L. Chacón, D.A. Knoll, J.M. Finn, *J. Comput. Phys.*, **178**, 15-36 (2002).
- [4] L. Chacón, D.A. Knoll, *J. Comput. Phys.*, **188** (2), 573-592 (2003).
- [5] L. Chacón, D.A. Knoll, *Proc. 33<sup>rd</sup> EPS Conference on Plasma Physics*, June 19-23, 2006, Rome, Italy.
- [6] L. Chacón, *Comput. Phys. Comm.*, **163** (3), 143-171 (2004).
- [7] B. Philip, M. Pernice, and L. Chacón, *Lecture Notes in Computational Science and Engineering*, accepted (2007).

#### Funding Acknowledgments

- Department of Energy, Office of Science, Office of Advanced Scientific Computing Research

# An Optimal Robust Equidistribution Method for 2D Grid Generation Based on Monge-Kantorovich Optimization

Gian L. Delzanno, Luis Chacón, John M. Finn, Yeo-Jin Chung, Giovanni Lapenta, T-15

**E**quidistribution of a positive monitor function (a density) is a central concept in grid adaptation. An example of its utility is in equidistribution of truncation errors in a discretization scheme, leading to minimization of the total error.

We have developed a new equidistribution method for 2D grid adaptation, based on Monge-Kantorovich optimization. The method is based on a rigorous variational principle, in which the  $L_2$  norm of the grid displacement is minimized, constrained *locally* to produce a prescribed positive cell volume distribution. This procedure minimizes the grid velocity in a time-stepping context, and avoids the central problem of Lagrangian methods: grid tangling. Our method is in contrast with commonly used variational grid adaptation schemes, in which a linear combination of cost integrals measuring different grid properties is used. In these methods, the scheme is a compromise between these different measures, and neither prevention of grid tangling nor exact equidistribution is achieved.

Our method involves finding a new grid  $x'(\xi)$ , ( $\xi$  is a logical rectangular grid) with density  $\rho'(x')$ , in terms of the old grid  $x$  with density  $\rho(x)$ . Optimization leads to  $x' = x + \nabla\Phi$  with

$$\nabla^2\Phi + H[\Phi] = \frac{\rho(x)}{\rho'(x')} - 1, \quad (1)$$

where  $H$  is the Hessian of  $\Phi$ . This is the Monge-Ampère equation, a single, nonlinear elliptic scalar equation with no free parameters and with well-known existence and uniqueness properties. Once the solution is found, the adapted grid should not fold since volumes are prescribed to be positive, and displacement of grid points is minimized. We solve this equation numerically with a Jacobian-free Newton-Krylov method. The ellipticity property of Eq. (1) allows multigrid

preconditioning techniques to be used effectively. The resulting numerical method is algorithmically scalable and uses a negligible amount of computational time and storage.

We have also shown that our variational principle maximizes grid smoothness (minimizes grid distortion), as measured by the  $L_2$  norm of the trace of the metric tensor. In Figs. 1-3, we show the resulting grid for two challenging cases of densities to be equidistributed. We plan to extend this grid generation formulation based on Monge-Kantorovich optimization to three dimensions.

**For more information contact John M. Finn at [finn@lanl.gov](mailto:finn@lanl.gov).**

## Funding Acknowledgments

- Department of Energy, Office of Science, Office of Advanced Scientific Computing Research
- Department of Energy, Office of Science, Office of Fusion Energy Sciences




Fig. 1. The optimal grid that equidistributes a density  $\rho'(x')$  (with  $\rho(x) = 1$ ) simulating spiraling arms of vorticity in the nonlinear Kelvin-Helmholtz instability, in which Lagrangian codes can develop grid tangling.




Fig. 2. Grid lines for a flower image.




Fig. 3. Gridding up the ubiquitous Lena.




# High-order Mimetic Finite Difference Methods on Arbitrary Meshes

Konstantin Lipnikov, T-7; Vitaliy Gyrya, Pennsylvania State Univ.

The mimetic finite difference (MFD) methods mimic important properties of physical and mathematical models. As a result, conservation laws, solution symmetries, and the fundamental identities of the vector and tensor calculus are held for discrete models. The existing MFD methods for solving diffusion-type problems on arbitrary meshes are *second-order* accurate for the conservative variable (temperature, pressure, energy, etc.) and only *first-order* accurate for its flux. In many physical simulations such as reactive transport in porous media and compressible flows, the flux accuracy makes significant impact on evolution of conservative quantities. We developed new high-order MFD methods that are *second-order* accurate for both the conservative variable and its flux [1]. These methods are well suited for simulations on arbitrary polygonal meshes.

Modelling with arbitrary polygonal meshes has a number of advantages. Such meshes allow us to describe accurately small, detailed structures such as tilted layers, irregular inclusions, rugged surfaces, and interfaces. The polygonal meshes cover the plane more efficiently than triangular meshes, and eventually reduce the number of discrete unknowns without loss of accuracy. The locally refined meshes, used to improve resolution in regions of interest, are particular examples of polygonal meshes with degenerate elements.

There are a few fundamentally different approaches for increasing the accuracy of discretization methods. Finite volume and finite difference methods increase stencils of discrete operators, which impose severe restrictions on mesh smoothness. These methods are usually applied on smooth meshes and lose accuracy on rough ones. The finite element and spectral element methods increase the number of unknowns inside each element but impose severe restrictions on the shape of admissible mesh elements. To develop new high-order MFD methods, we blend ideas of the finite element [2] and the low-order MFD methods [3]. Sample meshes for analysis are shown in Fig. 1; Fig. 2 shows the convergence rates for the low-order and high-order methods, and Fig. 3 illustrates the degrees of freedom in both methods.



*Fig. 1. Sample meshes used in analysis. Both randomly perturbed (left) and polygonal (right) meshes are challenging tests for any discretization method. The new high-order MFD methods have similar approximation properties on both meshes.*

In our analysis we consider a stationary diffusion problem for the conservative variable  $u$  and its flux  $\vec{F}$ :

$$\operatorname{div} \vec{F} = q, \quad \vec{F} = -K \nabla u.$$

In the high-order MFD method, the scalar function  $u$  is represented by one unknown, its average value, in each mesh element. The flux  $\vec{F}$  (the vector function) is represented by two unknowns on each mesh edge, the number which is twice more than in the low-order MFD method.

Similar to the low-order MFD method, the key step in the high-order MFD method is the definition of the inner product in the space of discrete fluxes. This inner product can also be viewed as a quadrature rule for the integral of a dot-product of two continuous fluxes. Since the chosen degrees of freedom are normal fluxes on mesh edges, the construction of this inner product is a non-trivial task. Due to additivity of integration, this inner product can be defined independently on each mesh element. We developed two methods for building elemental inner products [1].




Fig. 2. Convergence rates for the low-order MFD (blue) and the new high-order MFD (red) methods on randomly perturbed quadrilateral meshes for a manufactured solution. Both methods are second-order accurate for the conservative variable  $u$ . The high-order MFD method is second-order accurate for the flux  $F$ , while the low-order MFD method is only first-order accurate.




Fig. 3. Schematic illustration of degrees of freedom in the low-order left, and high-order right, MFD methods. One arrow represents the average normal flux through the edge. Two close arrows represent the average normal flux (0-th moment) through the edge and its first moment. The solid circles represent degrees of freedom for the conservative variable.

The first method extends further the ideas of the low-order Kuznetsov-Repin finite element method [3]. We divide *virtually* each polygonal element into triangles and use the existing formula for exact integration of linear fluxes on a triangle [4]. The virtual triangular partition introduces additional flux unknowns on interior edges. Half of these unknowns, 0-th moments of the normal flux, are eliminated using the Kuznetsov-Repin approach. The remaining unknowns first-moments are eliminated by mimicking integral identities for particular spaces of vector functions. In the finite element community this technique is known as the static condensation. The method is useful for problems where the flux has to be recovered at some points inside a polygonal element.

The second method was inspired in part by the methods developed in [3]. Only boundary data (normals to polygon edges, length of edges, and quadrature rules for edge integrals) are used to build the elemental inner product. Since no auxiliary triangular partition is required, the proposed method can be easily applied on meshes with degenerate polygons, which appear in adaptive mesh refinement (AMR) methods, and nonconvex polygons, which appear in moving mesh methods.

For more information contact Konstantin Lipnikov at [lipnikov@lanl.gov](mailto:lipnikov@lanl.gov).

- [1] V. Gyrya and K. Lipnikov, "High-order mimetic finite difference method for diffusion problems on polygonal meshes," Los Alamos National Laboratory report LA-UR-07-7196 (2007).
- [2] F. Brezzi, J. Douglas, Jr., and L.D. Marini, *Numer. Math.*, **47** (2), 217-235 (1985).
- [3] F. Brezzi, K. Lipnikov, and V. Simoncini, *Math. Models Methods Appl. Sci.*, **18** (3), 26-278 (2003).
- [4] Y. Kuznetsov and S. Repin, *J. Numer. Anal. Math. Modelling*, **18** (3), 26-278 (2003).

#### Funding Acknowledgments

- Department of Energy, Office of Science, Office of Advanced Scientific Computing Research

# A Multilevel Multiscale Mimetic ( $M^3$ ) Method for Two-phase Flows in Porous Media

Konstantin Lipnikov, David Moulton, Daniil Svyatskiy, T-7

**F**low simulations in porous media involve a wide range of strongly coupled scales. The length scale of short and narrow channels is on the order of millimeters, while the size of a simulation domain may be several kilometers (the richest oil reservoir in Saudi Arabia, Ghawar, is 280 km x 30 km). The permeability of rock formations is highly heterogeneous and may span several orders of magnitude, from nearly impermeable barriers to high-permeable flow channels. For such complex systems fully resolved simulations become computationally intractable. To address this problem we developed the new multilevel multiscale mimetic ( $M^3$ ) method [1]. This method possesses several distinctive features that lead to more reliable, robust, and efficient simulations of subsurface flows:

**Upscaled model.** Using the same mathematical model with averaged parameters to perform simulations at a much coarser scale does not adequately capture the influence of the fine-scale structure. In contrast, the  $M^3$  method constructs a hierarchical sequence of coarse-scale models, which provides a framework to capture fine-scale effects more accurately.

**Multilevel hierarchy.** Many different model upscaling approaches have been proposed [2,3]. All of these methods, except the multilevel upscaling (MLUPS) method [3], consider a two-level structure: coarse- and fine-scale grids. Using a two-level structure most multiscale methods achieve a coarsening factor of approximately 10 in each coordinate direction, while the trends in fine-scale realizations of large reservoirs require a coarsening factor of 100 or more. Using the multilevel hierarchy of the  $M^3$  method we achieve large coarsening factors with small computational cost.

**Locally conservative velocity field.** A multilevel framework was developed in the MLUPS method, but this approach does not produce conservative velocity fields. This is a crucial requirement for modeling two-phase flows, as these are described by a coupled system of equations for pressure (elliptic) and water saturation (hyperbolic). The  $M^3$  method provides locally conservative velocity fields on all scales, which guarantees local mass conservation.

**Algebraic nature.** We merge two computational strategies that were never used for two-phase flow simulations. The first strategy is the algebraic coarsening developed by Y. Kuznetsov that reduces the degrees of freedom inside a coarse-grid cell [4]. The second is a *novel* approach to conservative coarsening of velocities on the edges of a coarse-grid cell. These complementary strategies ensure that the coarse-scale system has the same sparsity structure as the fine-scale system, which naturally leads to a multilevel algorithm. Due to its algebraic nature, the method can be adapted to other fine-scale discretizations, such as the mixed finite element and finite volume methods and can handle full permeability tensors and general polygonal meshes.

**Conservative coarsening.** The conservative coarsening procedure is defined by velocity coarsening parameters. These parameters play a critical role in the accuracy of the  $M^3$  method. We implement a black-box, problem-dependent, and computationally inexpensive strategy to estimate them. In most multiscale methods the specific parameters that define the coarsening procedure are computed at the initial time step, with high accuracy, and are not changed during the simulation. Our numerical experiments demonstrate that it is important to update these parameters in time, even with moderate accuracy. Thus we propose to update our velocity coarsening parameters a few times during the simulation (e.g., every 500 time steps) using an efficient algebraic multigrid algorithm with a modest convergence tolerance. With this update strategy the error in the  $M^3$  solution is comparable to the error in the fine-scale solution.

The  $M^3$  method has been applied to the upscaling benchmark from the 10th SPE comparative solution project. We simulated flow in the fluvial layer shown in Fig. 1. (left) with the five-spot well configuration shown in Fig. 2 (right). The permeability field has large channelized structures, which is a challenging problem for multiscale methods. To discretize this system in time we use the IMPES time discretization scheme (implicit pressure and explicit saturation). The saturation is updated using the Darcy velocities provided by the pressure solver.

The numerical results demonstrate that with a large coarsening factor, such as 30, the  $M^3$  solution is close to the fine-scale solution (see Fig. 2). In other numerical tests for larger problems we implement more aggressive coarsening with a factor of 64, and also observe good agreement with the fine-scale solution. The  $M^3$  method speeds up the pressure solver up to 80 times, and the overall simulation eight times, with respect to the fine-scale simulation.

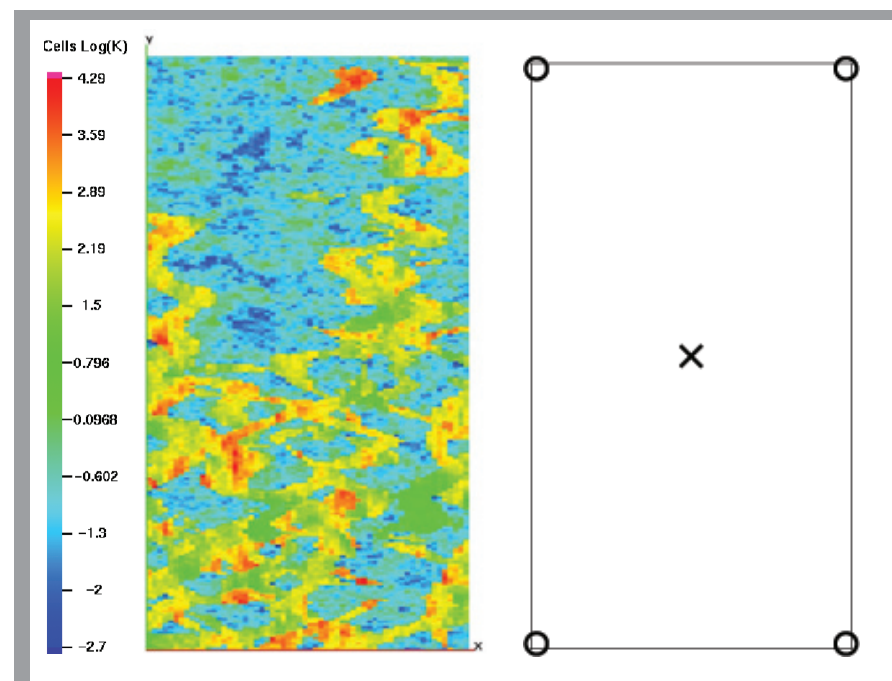


Fig. 1. The absolute permeability,  $K$ , of the 68th layer in an SPE benchmark model (left). Locations of the injector (x) and the producers (o) (right).

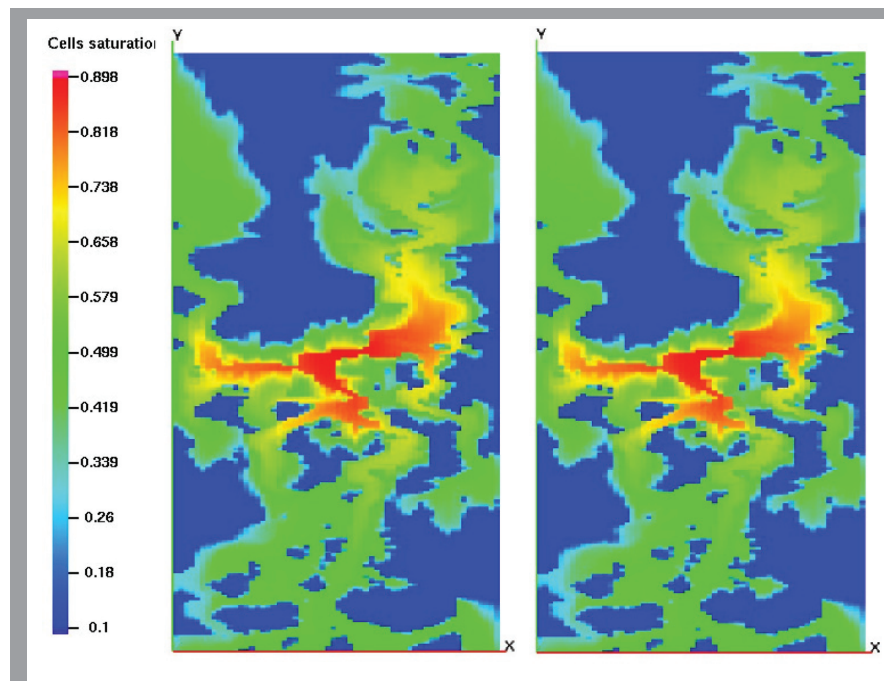


Fig. 2. The water saturation in the two-phase immiscible flow model after 880 days with injection of 200 ft<sup>3</sup> of water per day. The multiscale solution (right) preserves important features of the fine-scale solution (left). The pressure equation was solved on  $7 \times 2$  and  $220 \times 60$  meshes, respectively. The speedup of the pressure solver is 60 times.

For more information contact David Moulton at [moulton@lanl.gov](mailto:moulton@lanl.gov).

- [1] K. Lipnikov, D. Moulton, D. Svyatskiy. A multilevel multiscale mimetic ( $M^3$ ) method for two-phase flows in porous media. *Los Alamos National Laboratory Publication* 07-6160; Submitted to *J. Comput. Phys.* (2007).
- [2] V. Kippe, J. Aarnes, and K.-A. Lie, "A comparison of multiscale methods for elliptic problems in porous media flow," *Computat. Geosci.* (2007). To appear in special issue on multiscale methods.
- [3] S. MacLachlan and J. Moulton, *Water Resour. Res.*, **42**, W02418 (2006).
- [4] Y.A. Kuznetsov, *J. Numer. Math.*, **14** (4), 305-315 (2006).

#### Funding Acknowledgments

- Department of Energy, Office of Science, Office of Advanced Scientific Computing Research

# New Monotone Finite Volume Scheme for Anisotropic Diffusion

Konstantin Lipnikov, Mikhail Shashkov, Daniil Svyatskiy, T-7;  
Yuri Vassilevski, Russian Academy of Sciences

**P**redictive numerical simulations of subsurface processes require not only more sophisticated physical models but also more accurate and reliable discretization methods for these models. We have studied a new monotone finite volume scheme for diffusion problems with a heterogeneous anisotropic material tensor [1]. Examples of anisotropic diffusion include diffusion in geological formations, head conduction in structured materials and crystals, image processing, biological systems, and plasma physics. Development of a new discretization scheme should be based on the requirements motivated by both practical implementation and physical background. This scheme must meet the following conditions:

- be locally conservative;
- be monotone, i.e., preserve positivity of the differential solution;
- be applicable to unstructured, anisotropic, and severely distorted meshes;
- allow arbitrary diffusion tensors;
- result in sparse systems with a minimal number of non-zero entries;
- have higher than the first-order accuracy for smooth solutions.

The discretization methods used in existing simulations, such as the mixed finite element (MFE) method (Fig. 1), finite volume (FV) method, mimetic finite difference (MFD) method, and multipoint flux approximation (MPFA) method, satisfy most of these requirements except the monotonicity. They fail to preserve positivity of a continuum solution when the diffusion tensor is heterogeneous and anisotropic or the computational mesh is strongly perturbed. For instance, in simulations of a subsurface transport, a negative discrete solution of the pressure equation implies nonphysical Darcy velocities and hence wrong prediction of a contaminant transport.

Recently a few nonlinear monotone schemes have been suggested [2,3]. We studied schemes based on the nonlinear flux formula proposed in [3]. We rectified the LePotier's scheme for the case of unstructured triangulations and full diffusion tensors by giving correct positions of reference points. To improve robustness

of the scheme, we proposed an alternative interpolation technique [4]. We gave the first proof of scheme monotonicity for the steady diffusion equation. We studied numerically important features of the scheme such as violation of the discrete maximum principle and impact of the diffusion anisotropy on the scheme convergence. We extended the scheme to shape regular polygonal meshes and heterogeneous isotropic diffusion tensors.

The mixed form of the diffusion equation includes the mass conservation equation and the constitutive equation:

$$\begin{aligned}\operatorname{div} \mathbf{q} &= Q, \\ \mathbf{q} &= -\mathbb{D} \operatorname{grad} C,\end{aligned}$$

where  $\mathbb{D}$  is the diffusion tensor,  $Q$  is the source term, and  $\mathbf{q}$  is the flux of concentration  $C$ .

All the methods mentioned above use the same discretization of the mass conservation equation and differ by their approximation of the flux (constitutive) equation. In the nonlinear finite volume scheme a *reference* point is defined for each mesh cell  $T$  to approximate the concentration  $C$ . The position of the reference point depends on the geometry of  $T$  and value of the diffusion tensor. For isotropic diffusion tensors and triangular cell  $T$ , the center of the inscribed circle is the acceptable position for the reference point.

The flux  $\mathbf{q}$  is approximated at the middle of each mesh edge using a weighted difference of concentrations in two neighboring cells. Nonlinearity comes from the fact that these weights depend on a concentration at the edge vertices. To approximate solution at a mesh vertex, a linear interpolation method has been proposed in [3]. This method uses the three closest reference points, which form a triangle containing the vertex. We found out that this method is not robust for problems with strong anisotropy and sharp gradients. We proposed the *inverse distance weighting* interpolation method [4] for such a problem type. This method uses values at all reference points from the closest neighborhood of the vertex. Numerical experiments show that the new method is more stable for highly anisotropic problems.

The nonlinear finite volume method results in a sparse system whose dimension is equal to the number of mesh cells  $T$ . For triangular meshes, the matrix of this system has at most four nonzero elements in each row. To solve the nonlinear algebraic problem we use the Picard iterative method, which guarantees monotonicity of the discrete solution for all iterative steps.






Fig. 1. Mixed finite element method,  $C_{\min}^h = -1.7$




Fig. 2. Nonlinear finite volume method,  $C_{\min}^h = 0$

Profile of solution  $C^h(x,y)$  on the distorted triangular grid. Domain: unit square with the hole in the center. Problem: diffusion equation with highly anisotropic tensor. Ratio of tensor's eigen-values is  $10^3$ . Tensor is rotated with respect to coordinate axes on  $60^\circ$  clockwise.  $C^h = 2$  on the hole,  $C^h = 0$  on the boundary of unit square. Analytical solution satisfies maximum principle,  $0 \leq C(x,y) \leq 2$ . The MFE method produces nonphysical solution with strong negative values.

The computational results demonstrate the flexibility and accuracy of the scheme [2]. For sufficiently smooth solutions, we achieve the second-order convergence for concentration  $C$  and at least the first-order for flux  $\mathbf{q}$  in a mesh-dependent  $L_2$ -norm. For nonsmooth, highly anisotropic solutions we observe at least the first-order convergence for both unknowns.

For more information contact Daniil Svyatskiy at [dasvyat@lanl.gov](mailto:dasvyat@lanl.gov).

- [1] K. Lipnikov, M. Shashkov, D. Svyatskiy, Y. Vassilevski. "Monotone finite volume schemes for diffusion equations on unstructured triangular and shape-regular polygonal meshes." *J. Comp. Phys.* **227** (1), 492–512 (2007).
- [2] E. Burman and A. Ern, *Comptes Rendus Mathématique*, **338** (8) 641–646 (2004).
- [3] C. LePotier, *C. C. Acad. Sci. Paris, Ser. I* 341, 787–792 (2005).
- [4] D. Shepard, *Proceedings of the 23rd ACM National Conference*, 517–524 (1968).

#### Funding Acknowledgments

- Department of Energy, Office of Science, Office of Advanced Scientific Computing Research

# Predicting the Small Scales of $\alpha$ -Models for Turbulence

Susan Kurien, T-7; Evelyn Lunasin, Univ. of California, Irvine

The use of numerical models for the Navier-Stokes equations of fluid dynamics is an established practice in the study of turbulent flows. The calculation of flow from a model, instead of the underlying Navier-Stokes equations, allows for the use of less computing resource for a given flow. The so-called ' $\alpha$ -models' for turbulence typically use a smoothed velocity field to transport a rough velocity field. The smoothing is performed over a length scale  $\alpha$ . The expectation is that the smoothed field recovers the large-scale (greater than  $\alpha$ ) statistical properties, such as the energy spectrum, of the 'true' turbulent flow, which is not smooth. However, for scales smaller than  $\alpha$  the statistics are not easily deduced analytically since there are two participating 'velocities' that have different characteristic timescales and presumably different dynamics. In this work we use numerical simulations to arrive at an empirical hypothesis as to how one can predict the scaling of the energy spectrum of an  $\alpha$  model. We have thus deduced that the ambiguity in choice of dynamical timescales in  $\alpha$  models, due to the presence of two velocities, disappears when one simply considers the relevant conserved quantity and uses its timescale to govern the evolution of all other quantities of interest, including the energy.

Suppose we denote the smoothed velocity by  $u$ , and the rough (unsmoothed) velocity by  $v$ , is it  $u$ ,  $v$ , or some combination of the two that determines the statistics of scales smaller than  $\alpha$ ? In this work we examine this question numerically using two different  $\alpha$ -models, the Navier-Stokes- $\alpha$  (NS- $\alpha$ ) model and the Leray- $\alpha$  model. We study the energy spectra of these two models in the 2D case in the limit as  $\alpha \rightarrow \infty$ . We will call the models respectively NS- $\infty$  and Leray- $\infty$ . In this limit, *all* scales are smaller than  $\alpha$ , thus maximally extending the regime of interest. The power-law scaling of the energy spectrum may be conveniently related to the governing dynamics using simple dimensional analysis. By measuring the scaling exponent of the energy spectrum, we can thus infer which quantity governs the development of the sub- $\alpha$  scales. Based on the results of our simulations, we present the hypothesis that the statistics of the small scales of any  $\alpha$  model must be determined by the dominant (conserved) quantity that is transferred downscale in the regime of scales smaller than  $\alpha$ . The two models we study exhibit different




Fig. 1. Black curve – energy spectrum of the  $\alpha = 0$  (Navier-Stokes) equations, showing close to  $k^{-3}$  scaling in the enstrophy cascade range  $5 < k < 50$ , as expected. Red curve – energy spectrum of the NS- $\alpha$  model simulation for  $\alpha \rightarrow \infty$ , scales as  $k^{-7}$ , consistent with dynamics governed by the flux of the conserved enstrophy in this regime. Inset – compensated spectra.




Fig. 2. Energy spectrum for a  $4096^2$  simulation of the Leray- $\alpha$  model for  $\alpha \rightarrow \infty$ . The scaling is  $k^{-5}$  in the wavenumber region  $5 < k < 50$ , consistent with dynamics governed by the flux of the conserved enstrophy, different from that of Navier-Stokes- $\alpha$ , in this regime.



conserved integrals, and measurement of the power-scaling of their spectra confirms our hypothesis. This report is a summary of work published in [3,4].

The simulations were performed at increasing resolutions up to a maximum of  $4096^2$  data points. The conserved quantity for NS- $\alpha$  is the integral (or total) of the quantity  $\frac{1}{2} |\nabla \times v|^2$  known as the *enstrophy*. The scaling exponent of the energy spectrum is  $k^{-7}$  as shown in Fig. 1. The exponent -7 may be shown to be the only possibility if the dynamics is governed by the characteristic timescale of the flux of the conserved quantity. The conserved quantity for Leray- $\alpha$  is the integral of the quantity  $\frac{1}{2} |\nabla \times u| |\nabla \times v|$ . Fig. 2 shows that the spectrum for this model scales as  $k^{-5}$ , which may be shown to be the only possibility for a dynamics governed by the timescale for the flux of the conserved quantity.

Thus, the results from resolved simulations of two different  $\alpha$ -models, with different conserved quantities, support our hypothesis that the scaling of the spectrum should be governed by the characteristic timescale for the flux of the conserved quantity for the particular  $\alpha$  model.

**For more information contact Susan Kurien at [skurien@lanl.gov](mailto:skurien@lanl.gov).**

- [1] C. Foias, D.D. Holm, and E.S. Titi, *Physica D*, **152–153**, 505–519 (2001).
- [2] C. Foias, O. Manley, R. Rosa, and R. Temam, *Navier-Stokes Equations and Turbulence*, Cambridge University Press (2001).
- [3] E. Lunasin, S. Jurien, M. Taylor, and E.S. Titi, *J. Turbulence*, **8**, 751-778 (2007).
- [4] E. Lunasin, E.S. Titi, and S. Kurien, to appear in *J. Phys. A* (2008).

#### **Funding Acknowledgments**

- Department of Energy, Office of Science, Office of Advanced Scientific Computing Research

# Understanding the Small Scales in Rotating Stratified Flow

Susan Kurien, T-7; Beth Wingate, CCS-2; Leslie Smith, Univ. of Wisconsin

We present two significant new results from studies of the multiscale dynamics of rotating and stratified flows: 1) new parameterization of turbulence in the limit of strong rotation and stratification, and 2) a systematic way to improve upon the quasi-geostrophic (QG) approximation. The QG approximation is commonly used to describe mesoscales in the atmosphere and oceans but cannot capture coupling between wave motions and coherent structures (jets, fronts, vortices, and layers) on long time scales. Our work highlights the importance of proper parameterization of fast dynamics for accurate simulation of long-time weather prediction, ocean circulation, and climate change.

**Parameterization of the small scales in rotating and stratified flows.** An important goal for realistic computations of multiscale systems such as the ocean, atmosphere, or climate is accurate parameterization of complex physics in terms of simplified representations; for example, linear equations with coefficients determined from theory, experiments, and/or observations. Such parameterizations can then be incorporated into models. For example, in meteorological models, large-scale weather patterns evolve based on parameterization of small-scale turbulence.


A common approximation for rotating stratified flows is the so-called QG approximation. The two- and three-dimensional quasi-geostrophic (2DQG and 3DQG) equations can be derived by filtering all wave motions from their 'parent' equations, respectively the rotating shallow water (RSW) equations and the Boussinesq equations. The 2DQG and 3DQG equations provide a simplified framework for the study of large-scale, slowly varying atmospheric and oceanic dynamics in isolation from small-scale, rapidly varying fluid motions. Such study is often convenient but oversimplified because real physical systems such as the ocean and atmosphere exhibit strong coupling between the fast waves and the coherent large scales (jets, fronts, vortices, and layers) over long time scales.

We obtain a new theoretical prediction for the distribution of kinetic and potential energy in the turbulent small scales of rapidly rotating and stratified flow, without assuming quasigeostrophy. We find that kinetic energy is suppressed in the small, horizontal scales, while potential energy is suppressed in the small, vertical

scales. These constraints arise from studying the statistics of an important conserved quantity known as the potential enstrophy. Our results provide proper parameterization of small scales that are too expensive to calculate explicitly in realistic simulations of ocean/atmosphere dynamics.

**Capturing cyclone-anticyclone asymmetry.** Neither the 2DQG nor 3DQG equations can produce strong asymmetry between cyclones and anticyclones as is ubiquitous in geophysical flows; for example, the long-lived anticyclone known as Jupiter's Red Spot, the predominance of cyclones in the mid-latitude North Atlantic, and the fact that hurricanes are always cyclonic.

We have developed a novel approach to derive a hierarchy of models that successively include more and more effects of the wave motions inherent in the full RSW and Boussinesq systems. In sharp contrast to most previous efforts to improve upon QG models, our technique does not rely on a small parameter. This means that our new 'corrections' to the 2DQG and 3DQG model equations need not be small and are likely to capture essential physics over a wide range of rotation rates and stratification strengths.



*Fig. 1. Vorticity skewness vs time during decay from random initial conditions: full RSW (red); 2DQG (blue); new PPG model including QG interactions and all interactions involving one gravity wave (black). Negative skewness indicates the predominance of anticyclonic vortices.*

To illustrate the power of our new models, we compare decay dynamics of the full RSW equations, the 2DQG model equation filtering all wave motions, and a new model denoted PPG. The PPG model improves upon 2DQG by adding all nonlinear interactions involving one gravity wave. Figure 1 shows the evolution of the vorticity skewness during decay from random initial conditions, where negative skewness indicates the formation of anticyclonic vortices. One can see that the full RSW and PPG equations quickly develop a negative skewness, whereas the skewness associated with the 2DQG models remains close to zero, indicating roughly equal numbers of cyclones and anticyclones.

Our approach is based on a Fourier (wave) representation of the original RSW and Boussinesq equations, for which there are 27 types of nonlinear interactions (some redundant). Retaining a larger subset of interactions results in a model with more complexity. In physical space, the new models can be used with realistic boundary conditions, and they naturally conserve energy. The clear advantage over previous models is lack of dependence on a small parameter.

**For more information contact Susan Kurien at [skurien@lanl.gov](mailto:skurien@lanl.gov).**

[1] S. Kurien, B. Wingate and M. Taylor, "Anisotropic small-scale constraints on energy in rotation stratified turbulence," submitted to *Phys. Rev. Letts.* (2007).

[2] M. Remmel and L. M. Smith, "New intermediate models for rotating shallow water and an investigation of the preference for anticyclones," in preparation for *J. Fluid Mech* (2007).

### **Funding Acknowledgments**

- Department of Energy, Office of Science, Office of Advanced Scientific Computing Research
- Los Alamos National Laboratory Directed Research and Development Program

# Modeling Multiscale-multiphase-multicomponent Subsurface Reactive Flows Using Advanced Computing

Peter Lichtner, Chuan Lu, EES-6; David Moulton, Bobby Philip, T-7

**I**t is becoming increasingly clear that the ability to model multiscale subsurface processes is essential for obtaining an accurate predictive capability of contaminant transport [1]. Predictive modeling of reactive flows is a daunting task because of the wide range of spatial scales involved—from the pore to the field scale—ranging over more than six orders of magnitude, and the wide range of time scales involved—from seconds or less to millions of years. However, even with ultrascale (petascale and beyond) computing facilities, large 3D field scale continuum models employing billions of nodes on uniform grids can only resolve features on the order of tens of meters and cannot capture phenomena at much smaller scales on the order of centimeters to millimeters or less. Heterogeneity, multiphase interfacial effects, and multicomponent geochemical reactions add further complexity to the system requiring advances in modeling capabilities.

To meet these challenges, algorithms scaleable to large numbers of processors (>10,000) are needed to provide efficient simulation on ultrascale computing platforms. Adaptive mesh refinement (AMR) will be essential to limit the number of nodes needed to adequately represent the spatial variability of the system. And subgrid scale models will be essential to capture multicontinuum effects involving fracture-rock matrix interaction, preferential flow paths, and other multiscale phenomena. Implicit in applications of AMR to heterogeneous media is the need for upscaling (e.g., permeability tensor) to model reactive flows at different resolutions.

These enhanced modeling capabilities will be used to improve our understanding of radionuclide migration at the DOE Hanford facility, where sub-millimeter-scale mass transfer effects have thwarted attempts at remediation efforts, and modeling sequestration of CO<sub>2</sub> in deep geologic formations, where resolving density-driven fingering patterns is necessary to accurately describe the rate of dissipation of the CO<sub>2</sub> plume. The focus of this highlight is on the challenges facing modeling CO<sub>2</sub> sequestration in geologic formations.

**CO<sub>2</sub> Sequestration in Saline Aquifers.** CO<sub>2</sub> sequestration (capture, separation, and long-term storage) in various geologic media including depleted oil reservoirs,

saline aquifers, and oceanic sediments is being considered as a possible solution to reduce greenhouse gas emissions. Sequestration in subsurface geologic formations containing saline aquifers could provide permanent storage and thereby help mitigate global climate change. Saline aquifers have an estimated world-wide storage capacity for CO<sub>2</sub> of 100-10,000 Gt CO<sub>2</sub> [2].

Injection of supercritical CO<sub>2</sub> into deep underground reservoirs introduces unique computational challenges. At appropriate *p-T* conditions, supercritical CO<sub>2</sub> is buoyant because of its lower density compared with the brine, and rises, eventually becoming trapped by a caprock and spreading laterally as it gradually dissolves into the surrounding brine. As the CO<sub>2</sub> dissolves into the brine, however, the brine becomes heavier and begins to sink, resulting in density-driven convective instabilities leading in the formation of CO<sub>2</sub>-concentrated brine fingers protruding downward. Convective mixing can result in much more rapid dissipation of the supercritical CO<sub>2</sub> plume compared with diffusive processes alone. Also important is the areal extent of the plume, which can increase the accidental release of CO<sub>2</sub> to the surface through abandoned bore holes and faults. Competition between the rate of spreading of the plume and the rate at which it dissolves into the brine controls the extent of spreading, and this competition is a complex multiscale process that is difficult to model.

**Massively Parallel Multiphase Code PFLOTTRAN.** To investigate the role of convective mixing numerically, the massively parallel computer code PFLOTTRAN is used to model multiphase-multicomponent processes involved in sequestration [3]. PFLOTTRAN provides for phases of CO<sub>2</sub> and brine under nonisothermal conditions. Phase transformations are described using a variable switching approach as the thermodynamic state of a grid block changes from brine, to two phases: CO<sub>2</sub>-brine, to pure CO<sub>2</sub>.

The role of fingering on dissipation of a CO<sub>2</sub> plume is investigated numerically by injecting supercritical CO<sub>2</sub> into a highly permeable sandstone formation. As demonstrated in Fig. 1, the evolving profile of dissolved CO<sub>2</sub> is highly dependent on resolution of the fingers produced following the injection process. Fingering will also be important in determining chemical interactions between between reservoir minerals and the CO<sub>2</sub>-enhanced brine.

**For more information contact Peter Lichtner at [lichtner@lanl.gov](mailto:lichtner@lanl.gov).**

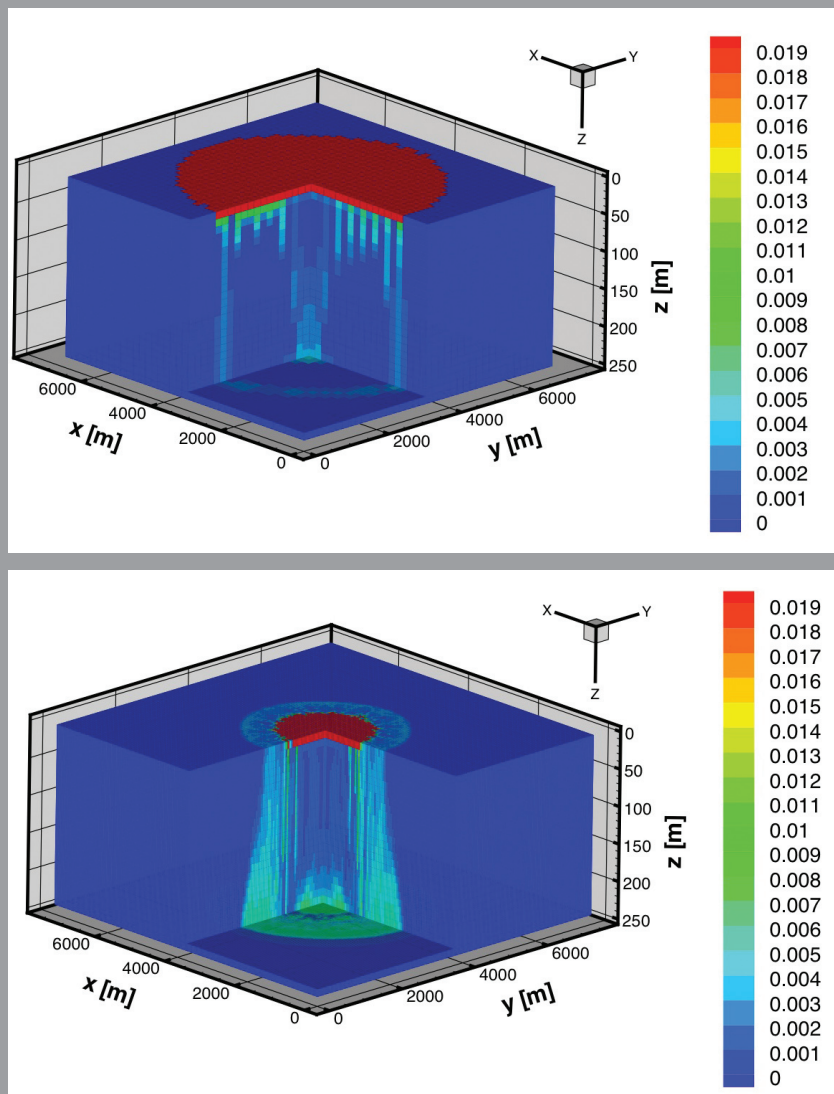



Fig. 1. Simulations carried out using PFLOTRAN showing dissolved  $\text{CO}_2$  at an elapsed time of 300 years illustrating the dependence of fingering on coarse (top) and fine (lower) grid resolutions. The coarse grid consists of  $40 \times 40 \times 25$  nodes with spacing  $D_x = D_y = 175$  m and  $D_z = 10$  m. The fine grid is refined in the  $x$  and  $y$  directions by a factor of four. An isotropic permeability of  $2 \times 10^{-12} \text{ m}^2$  is used with a porosity of 15% typical of sandstone. The computational domain is 250 m thick and  $7 \times 7$  km in lateral extent.  $\text{CO}_2$  is injected at a rate of 1 MT/y for 20 years at a depth of 50 m below the top of the domain, corresponding to roughly 75% of the  $\text{CO}_2$  produced by a 1000 MW gas-fired power plant in 20 years. No-flow boundary conditions are imposed at the top and bottom and front and back of the domain, and with constant pressure at the left and right sides.

- [1] Basic research needs for geosciences: facilitating 21 st century energy systems. Workshop sponsored by the U.S. Department of Energy, Office of Basic Energy Sciences, Feb 2007. (<http://www.sc.doe.gov/bes/reports/list.html>).
- [2] Ibid.
- [3] C. Lu, and P. Lichtner, *Journal of Physics Conference Series*, **78**, 012042 doi:10.1088/1742-6596/78/1/012042 (2007).

#### Funding Acknowledgments

- Department of Energy, Office of Science, Scientific Discovery through Advanced Computing Program
- Department of Energy, Office of Science, Innovative and Novel Computational Impact on Theory and Experiment Program







# Astrophysics and Cosmology

Astrophysics and cosmology have a long history at LANL. In fact, astrophysicists have been involved in the nuclear weapons program since its inception in 1943. More recently, astrophysicists and cosmologists have been instrumental in nonproliferation as well as other observational programs associated with the mission of the Laboratory.

In this section, there are four papers in astrophysics and cosmology with funding originating in “work for others.” The first paper discusses progenitors of gamma-ray bursts. Gamma-ray bursts were discovered by LANL scientists in the 1960s. They are the most luminous electromagnetic events in the universe since the Big Bang. The project focus is on the progenitors of long-duration bursts.

The second article in this group deals with the analysis of the huge astronomical data sets that are expected to become available in the next decade. These enormous sets of data, from large astronomical observation projects, will challenge even the fastest computers and the best algorithms for analyzing such data. This article describes strategies for dealing with the storage, analysis, and management of this data.

The third article describes state-of-the-art research in dark energy. Since the 1970s the presence of dark matter in the universe has been acknowledged from measurements on the speeds of stars in galaxies and

other observations. Only in the last decade has the presence an enormous amount of dark energy become a topic of research. This article summarizes some of the work being done in this area at LANL.

The final article in the series on astrophysics and cosmology deals with astrophysical explosions, gamma-ray bursts, and nucleosynthesis. This work is part of a project funded by the Department of Energy that involves researchers from Los Alamos and the University of California at Santa Cruz. The article describes how the researchers use supercomputers to simulate these most violent explosions.

# Studying Gamma-ray Burst Progenitors under the NASA Swift Guest Investigator Program

Chris L. Fryer, Gabriel Rockefeller, Patrick A. Young, CCS-2

**W**e have been funded to study gamma-ray burst (GRB) progenitors for the past 3 years. Each year, the collaborators, generally postdocs who work with Chris Fryer, change, as does the exact emphasis of the project when these postdocs take positions elsewhere. Last year's project was led by Chris Fryer, Patrick Young, and Gabriel Rockefeller. Gabriel is continuing a Director's postdoc here at LANL under Aimee Hungerford, CCS-2.

In this year, we have made considerable progress understanding the progenitors of long-duration GRB. We finished five papers [1,2,3,4,5] since November 2006 studying GRB progenitors, with six additional papers [6,7,8,9,10,11] studying physics relevant to GRB progenitors; studies either of massive-star collapse, supernova explosion or fallback, or binary properties. We led a comprehensive review of both the theory and observations that can be used to constrain the progenitors of GRBs [2]. In this work, which required the agreement of 17 scientists in the field, we surmised that binary GRB progenitors match the current set of data much better than single star models, even the rotationally induced mixing models [12]. At this point, we are reasonably convinced that the primary progenitor for long-duration GRBs requires a binary, and we will continue to focus our work on binary progenitors. This major review led to two very important lines of new research: the difference between black hole formation scenarios and the importance of understanding population synthesis uncertainties to make detailed observational comparisons.


First, our detailed study of the observational constraints led to a more focused study of long-duration GRB progenitors [3], discovering that the observations already can help us distinguish between systems that form black holes after the launch of a shock, fallback black holes, versus those systems forming black holes directly. Theory argued that some fallback black hole systems would not form GRBs if the fallback accretion rate is too low, but no firm predictions have been made, and some GRB scientists have put forth models with very low accretion rates.

Observations may be able to shed some light on this subject. Using metallicity and red-shift observations of GRBs, we found that the number of GRBs from fallback

black holes is limited to a few percent of the currently predicted fallback black-hole rate, suggesting that most fallback systems do not form the GRBs in our observed sample. Our first study in this topic was fairly simple, but we do not believe the main result discussed here will change with new observations. One of the projects for this year is to determine other observational constraints distinguishing between fallback and direct-collapse GRBs to confirm or deny our claim. One such effort stems from our realization that fallback black hole GRBs have different nucleosynthetic yields and associated supernova light curves than direct collapse black holes [1,4]. These discoveries also led us to develop new studies of collapsar progenitors and supernova shock propagation to better understand fallback [6,7,8,9].

Our review also helped us focus on some of the crucial uncertainties in studying binary populations that must be understood before making more direct constraints to the observations. We have studied in more detail the role of tidal effects in binaries [5] and wrapped up some observational studies of initial parameters, mass, and separation distributions for massive binaries [10,11]. We believe that if we can better understand the progenitors of normal type Ib/c supernovae, which our review [2] argued were also produced in binaries, we may gain some insight into the progenitors of GRBs. We have also done some type Ib/c population synthesis calculations [2,11].

We are currently on cycle 3 Swift funding. This funding covered work by Patrick Young examining aspects of GRB progenitors. It will also fund the work by Gabriel Rockefeller studying how angular momentum affects the fate of stellar




*Fig. 1. Number of GRBs from a combination of all the direct-collapse black holes with 2% of the total number of fallback black hole systems. Contours are plotted for bins with 20, 40, 60, and 80% of the peak number. If fallback black holes are dominant GRB progenitors, the average metallicity would be higher and the redshift lower. The rotationally induced mixing model can not explain GRBs with metallicities above 0.1 solar.*


collapse. Preliminary work (Fig. 2) on this showed that if the specific angular momentum in the material is too low, even though a disk forms, it is still unstable and unlikely to make the focused jets seen in observations. The nominal average angular momentum required to make a stable disk is  $10^{17} \text{cm}^2 \text{s}^{-1}$ , higher than that assumed in many population studies. This work has yet to be published, and we will spend some time with our current funding finishing this paper. We will also finish incorporating magnetohydrodynamics capabilities into our SNSPH [13] code to get an estimate of the role of magnetic fields in this collapse. Several other groups [14,15] have done some work in this area, and we will be able to compare our results with their work.

We will also spend part of the time of our current funding adapting the SNSPH code to better model the mass transfer in binaries. Steven Diehl, a postdoc at LANL, has already begun an effort to model low-mass binaries under an internal LANL effort led by F. Timmes to study hydrogen-deficient stars. Figure 3 shows some preliminary results from this effort, the common envelope evolution of the white dwarf (WD0137-049) brown dwarf binary [16,17]. As part of this effort, we have finished much of the verification and validation of this modified code, paying particular attention to the accuracy of the gravity routine and the ability for our scheme to model angular momentum conservation.

With the progress we have made, we expect to renew this NASA funding and continue to contribute LANL excellence in GRB progenitors.



**Fig. 2.** Plots looking down the angular momentum axis (z-axis) of two collapse calculations of a 60 solar mass star at 0.44s: one with an average specific angular momentum  $\sim 5 \times 10^{16} \text{cm}^2 \text{s}^{-1}$  (left), the other with twice that angular momentum (right). To get a stable accretion disk that can develop focused jets, we required this latter ( $10^{17} \text{cm}^2 \text{s}^{-1}$ ) high specific angular momentum.



**Fig. 3.** Simulation of the common envelope evolution of a brown dwarf within a Red Giant star using our SNSPH code [15,16].

For more information contact Chris L. Fryer at [fryer@lanl.gov](mailto:fryer@lanl.gov).

- [1] C.L. Fryer, P.A. Young, and A.L. Hungerford, *ApJ*, **650**, 1028 (2006).
- [2] C.L. Fryer, et al., accepted by *PASP*, astro-ph/0702338 (2007).
- [3] P.A. Young and C.L. Fryer, accepted by *ApJ*, astro-ph/0703373.
- [4] C.L. Fryer, A.L. Hungerford, and P.A. Young, *ApJ*, **662**, L55 (2007).
- [5] K. Belczynski, et al., *ApJ*, **664**, 986 (2007).
- [6] C.L. Fryer, *NewAR*, **50**, 492 (2006).
- [7] C.L. Fryer, A.L. Hungerford, and G. Rockefeller, *IJMPD*, **16**, 941 (2007).
- [8] C.L. Fryer, and P.A. Young, *ApJ*, **659**, 1438 (2007).
- [9] P.A. Young, and C.L. Fryer, *ApJ*, **664**, 1033 (2007).
- [10] D. Kiminki, et al., *ApJ*, **664**, 1102, (2007).
- [11] H.A. Kobulnicky, and C.L. Fryer accepted by *ApJ*, astro-ph/0605069.
- [12] S.-C. Yoon, and N. Langer, *A&A*, **460**, 199 (2006).
- [13] C.L. Fryer, G. Rockefeller, and M.S. Warren, *ApJ*, **643**, 292 (2006).
- [14] L. Dessart, A. Burrows, E. Livne, and C. Ott, astro-ph/0710.5789.
- [15] D. Proga, *RMxAC*, **27**, 47 (2007).
- [16] <http://www.eso.org/public/outreach/press-rel/pr-2006/pr-28-06.html>.
- [17] S. Diehl, C.L. Fryer, and F. Herwig astro-ph/0711.0322.

### Funding Acknowledgments

- National Aeronautics and Space Administration

## Astronomical Data Analysis with Commodity Components

Michael S. Warren, T-6; John Wofford, Columbia Univ.

**D**uring the next decade, large astronomical observing projects will generate more than 1000 times as much observational data as has been gathered in all of our history. Storage, analysis, and management of this information will require significant advances in computing technology. As an initial step in this process, we have developed and deployed an astronomical data system based on open-source software and commodity hardware with a storage capacity of 112 terabytes (TB) in immediately accessible disk arrays for a total cost of \$95,127. This approach is scalable to a petabyte of storage for less than \$1M. In the same way in which special-purpose telescopes are now required to obtain the best catalogs of objects in the sky, a focused effort involving state-of-the-art parallel computer hardware and software is required to analyze this data and model the evolution of the Universe, which led to the observed distribution of stars and galaxies.

Persistent data storage is a fundamental prerequisite for all information science and technology projects. The advent of commodity microprocessors with adequate floating-point performance and low-priced fast ethernet switches contributed to the emergence of Beowulf clusters in the mid-1990s, which revolutionized parallel computing at the departmental scale [1-3,4]. We are currently in the midst of a similar revolution in scalable data storage, due to the dramatic decline in the price of commodity disk drives and 10-gigabit (Gb) networking technologies.

Vast amounts of relatively inexpensive storage offer significant opportunities to develop new approaches to scientific problems. The cost for SATA disk storage is currently about \$0.35 per gigabyte for 1-TB drives. A single fault-tolerant RAID-6 node can store 14 TB in 3U of rack space for a total cost of about \$10 K (see Table 1). Used in a parallel cluster environment, a petabyte disk array with achievable read/write bandwidth that greatly exceeds available local and wide-area networking technology is possible.

The problem we address here is optimizing the price/capacity of information storage at the 100-terabyte-to-petabyte scale, while maintaining acceptable application performance within the domain of astronomical data processing applications. Implied in this price/capacity goal is minimizing overall costs over the

life-cycle of the system, including administration and maintenance. Acceptable application performance also includes minimizing the possibility of data loss and downtime due to hardware failures, while realizing that there are trade-offs in complexity and cost vs reliability.

The optimal way to move large files was with the *mpscp* software from Sandia National Laboratory, which, using four simultaneous TCP streams, obtained transfer rates of 380 MB/s between two nodes of the cluster, which is near the maximal rate at which a node can write data. We performed a CRC checksum on the 3.5 TB of Sloan Digital Sky Survey image data, consisting of 1.6 million files, in 2,364/s, for an overall processing rate of 686 files per/s, with an aggregate read bandwidth of 1.5 GB/s. Running *sextractor* to extract astronomical objects from each image file completed in 28686/s, for a processing rate of 56 images per/s, or equivalently, 172 megapixels/s. Performing a plane-to-plane reprojection using *mProjectPP* from Montage processed 112,000 files in 54,700/s, for an overall rate of 2 images per/s. An example reprojected image is shown in Fig. 2.

We executed a 32-processor parallel job on the eight nodes of the cluster using our N-body code with a test data set of 400 million particles. The actual write bandwidth to disk during the data dump was 2.6 GB/s. Restarting the code from a data dump required reading in the 12.8-GB/s file, which resulted in a read bandwidth of 1.8 GB/s. The overall floating point performance obtained by the cluster executing the N-body code was 97.8 Gflops. We created an animation



*Fig. 1. The deployed storage system, consisting of eight 3U storage nodes and an 8-port 10-gigE switch. The usable storage capacity of the system is 112 TB. The total cost of the cluster was \$95 K.*



of the evolution of the dark matter in the Universe, which required reading in every data dump produced by a simulation, and projecting the 3D mass density into a 2D image. For this analysis task, the cluster supported reading each 12.8 GB file in 7.5/s, for a read bandwidth of 1.7 GB/s. An example image from this process is shown in Fig. 3.

Reliability of the system has been excellent, with no disk failures in eight weeks of operation. This result is consistent with Google's recent statistics on using over 100,000 disk drives [2].

Our experience with the described hardware configuration has been positive. It provides a reasonable environment to work within, while providing immediately accessible disk storage at a cost that is almost within a factor of two of the cost of bare disk drives. The software environment should be familiar to any Linux user, and does not require any specialized system administration skills. In summary, we have demonstrated a flexible and usable storage system capable of storing 112 TB of information at a cost of \$95,127, or a price/capacity of \$850 per terabyte.

This project won the Storage Challenge Award at the annual SC2008 conference on high-performance computing, networking, and storage in Reno, NV.

**For more information contact Michael S. Warren at [msw@lanl.gov](mailto:msw@lanl.gov).**


- [1] D.J. Becker, et al., In Proceedings of the 1995 *International Conference on Parallel Processing (ICPP)*, 11-14 (1995).
- [2] E. Pinheiro, W.D. Weber, and L.A. Barroso, "Failure trends in large disk drive population," Proceedings of the 5th *USENIX Conference on File and Storage Technologies (FAST07)* (2007).
- [3] M.S. Warren, et al., Proceedings of the *International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA'97)*, 1372-1381 (1997).
- [4] M.S. Warren et al., "Pentium Pro inside: I. A teecode at 430 Gigafllops on ASCI Red, II. Price/performance of 50 Mflop on Loki and Hyglac," *Supercomputing '97*, Los Alamitos, IEEE Comp. Soc. (1997).

### Funding Acknowledgments

- National Aeronautics and Space Administration



*Fig. 2. An example image from the Sloan Digital Sky Survey that was reduced and reprojected using the system described here.*



*Fig. 3. An image representing the distribution of dark matter in the universe, which was produced from a 128 million particle N-body simulation using the system described here.*

## Going Nonlinear with Dark Energy

Salman Habib, T-8; Michael S. Warren, Darren Reed, T-6; Katrin Heitmann, ISR-1; Martin White, Lawrence Berkeley National Laboratory

The discovery and confirmation of cosmic acceleration is less than a decade old, yet it is already widely recognized as a harbinger of a revolution in fundamental physics and cosmology. The acceleration could be due to a mysterious dark energy with a peculiar equation of state or it could signal a breakdown of general relativity at very large scales. If the dark energy is literally constant, then its numerical value is ridiculously small compared with theoretical expectations. Although the theoretical existence of this problem has been appreciated since the 1930s, it was widely assumed that a fundamental symmetry was responsible for tuning the cosmological constant exactly to zero. Now, not only do we know this expectation to be false, but we also have to explain its current value—roughly twice that of the energy density in matter in the Universe. Aside from the cosmological constant, there exist dynamical models of dark energy termed quintessence, where the equation of state is time-varying. As long as we do not know whether the dark energy is constant or not, we also do not know what theoretical approach to take to attack the basic puzzle. Thus the very first step is to characterize the equation of state of dark energy as accurately as possible from observations. The challenge posed by next-generation observations has been laid out in the Department of Energy (DOE) Dark Energy Task Force Report [1]. From a theoretical and modeling perspective, the key requirement is a one to two orders-of-magnitude improvement in control of systematic effects.

Cosmological probes of dark energy rely on measuring the expansion history of the Universe and the growth rate of structure formation. For deployment in the near- and medium-term future, four techniques have been put forward. These are baryon acoustic oscillations, and cluster, supernova, and weak lensing cosmological surveys. Of the structure formation-based techniques, both weak lensing and cluster observations can directly measure the growth of structure as well as the expansion history, and hence can test whether general relativity is modified at large scales. The control of systematics in weak lensing is, however, much further advanced than for clusters. The aim of our project is to understand and control the main sources of theoretical errors underlying the baryonic oscillations and weak lensing-based dark energy research program by carrying out large-scale numerical simulations (Fig. 1) to refine and calibrate algorithms and analytic approximations, and potentially serve

as templates when the data become available. To support this project, a very significant allocation of institutional computing resources has been made available by LANL.

The source of baryon acoustic oscillations is the tight coupling of the baryon-photon fluid in the early Universe. During this phase, perturbations do not grow but propagate as sound waves. After recombination, the sound waves survive as imprints on the late-time distribution of matter at a scale of 150 Mpc, a fixed “cosmic ruler,” which in turn enables accurate measurements of the Hubble parameter and the angular diameter distance. Weak gravitational lensing refers to the distortion of background galaxy images due to the gravitational bending of light rays caused by foreground mass concentrations, the distortion being described in terms of image stretch (“shear”) and magnification (“convergence”). Measurements of shear and convergence directly map the distribution of matter in the Universe.

The first global aim of the project is to ensure that the numerical simulations to be carried out can provide a sufficiently complete physical description, yet one that is robust and possesses controlled errors and convergence properties—the hallmark of precision cosmology [2]. The number of simulations, the box sizes used, the initial condition accuracy, and the force and mass resolution needed are all important ingredients in determining the actual suite of simulations that will be performed. A key question is whether enough precision information can be gleaned from a relatively small number of accurate large-scale simulations. Recent investigations by members of the proposing team have made significant progress in this area [3,4]. An example of this approach is shown in Fig. 2. From a small 32-simulation sample, varying five cosmological parameters as well as the cosmological epoch, we can produce prediction results at the 1 percent accuracy at parameter values not included in the original simulation sample.

Another area that has been investigated is the precision of cosmological initial conditions and control of errors in the evolution of the density field on scales relevant to baryonic oscillation observations and weak lensing measurements. Our work has established that subpercent error initial conditions can be implemented and that error control in the evolution can be kept at this level down to length scales of the order of a few Mpc, which is sufficient for the analysis of current and near-term observations. A large suite of cosmological simulations is now being run with the aim of constructing a database for support of cosmological surveys. This database will be used by project members and will be made available to the wider dark energy community in the first half of 2008.






Fig. 1. The density field from a cosmological simulation showing the very large dynamic range in density that needs to be captured to make theoretical predictions for cosmological large-scale structure observations.




Fig. 2. Holdout error test showing accuracy of power spectrum emulator for weak lensing observations at three redshifts, based on the methodology of [2]. The dark gray zone contains 50 percent of the results, and the light gray, 90 percent. The control on accuracy is at the 1 percent level.

For more information contact Salman Habib at [habib@lanl.gov](mailto:habib@lanl.gov).

- [1] A. Albrecht et al., *astro-ph/0609591*.
- [2] K. Heitmann, et al., *Comp. Sci. Disc.* (invited) *arXiv:0706.1270 v1*.
- [3] K. Heitmann, et al., *Astrophys. J. Lett.* **646**, L1 (2006).
- [4] S. Habib et al., *Phys. Rev. D*, **76**, 083503 (2007).

#### Funding Acknowledgments

- Department of Energy, Office of Science, High Energy Physics Program

# Supernovae, Gamma-ray Bursts, and Nucleosynthesis

Alexander Heger, T-6

Our team of investigators, headed by principle investigator (PI) Stan Woosley at the Univ. of California, Santa Cruz, has begun the second phase of a DOE Scientific Discovery through Advanced Computing (SciDAC) project to understand stars and the explosive phenomena they produce, especially supernovae of all types, gamma-ray bursts (GRB), and x-ray bursts. We use supercomputer simulations to study the most violent explosions in the universe: supernovae, the powerful explosions of dying massive stars, and gamma-ray bursts, short and bright blasts of intense hard radiation first discovered by LANL scientists. By studying the evolution of massive stars and their explosion as supernovae and gamma-ray bursts, we gain new insights into how the “heavy” elements needed for life, such as oxygen and iron, are forged inside stars through the process of nucleosynthesis. When stars explode, the heavy elements created within them are ejected into space, eventually forming new stars and planets.

Models of these phenomena share a common need for nuclear reactions and radiation transport coupled to multidimensional fluid flow. Our principal goals are not only a better first-principles understanding of supernovae, GRBs, and nucleosynthesis, but also theoretical databases that will allow the more precise and reliable use of supernovae for cosmological distance determination. Our studies of nucleosynthesis in stars and supernovae will be the most complete in the world and will highlight nuclear uncertainties (especially in the *r*-process and *rp*-process) that will someday be elucidated by the next generation of nuclear physics experiments. In the next 5 years our program will be the only way to address significant gaps in our understanding of these objects and potential systematics in their use as cosmological probes. Our research will be able to directly influence the construction of these experiments, thereby ensuring that they are optimally designed to confront the greatest mystery in high-energy physics and astronomy today: the nature of dark energy.

The central issue in the formation of primordial stars (Population III stars) is their mass. The mass determines the nucleosynthetic yield, the nature of the supernova explosion, whether a black hole remnant is left, and whether the supernova leads to a GRB. We have shown that, in the absence of mass loss, primordial stars between 140 and 260 solar masses explode as pair-instability supernovae. Those higher in mass collapse to form black holes. Simulations show that Population



*Fig. 1. The illustration shows what the brightest supernova ever recorded, known as SN 2006gy, may have looked like. The fireworks-like material (white) shows a pulsational pair instability outburst from a very massive star. In a first supernova the star expelled tens of solar masses of gas that cooled down (red). As the material from the subsequent supernova explosion crashes into the lobes, it heats the gas in a shock front (green, blue, and yellow) and pushes it backward. (Credit: NASA/CXC/M.Weiss).*

III stars form from large condensations of gas and dark matter in which the Jeans mass, i.e., the critical mass for collapse of the gas cloud to a star, is on the order of 500 solar masses. The conclusion from both numerical and analytic studies is that the collapse of this gas does not result in fragmentation. As a result, the final mass of the star is set by feedback processes, the most important of which are photoionization and the associated radiation pressure. It is essential to simulate the formation of the first stars, including these feedback processes, to better constrain their initial masses.


Stars with helium core masses heavier than about 45 solar masses experience a dynamic instability after carbon burning. The conversion of energy into rest mass of electron-positron pairs at temperatures in excess of about  $10^9$  K robs the star of energy that would have gone to increasing the pressure, and as a result the core of the star implodes. This implosion continues until either explosive nuclear burning delivers enough energy to turn it around into an explosion, or a black

hole is formed. For nonrotating stars with helium core masses less than about 130 solar masses, an explosion is the outcome. Kinetic energies of up to  $10^{53}$  erg are possible, with some explosions producing as much as 50 M of iron (as radioactive  $^{56}\text{Ni}$ ). Successful 1D explosion models have been calculated for 40 years, and a recent survey [1] has elucidated the nucleosynthesis expected for the entire range of pair-instability supernova masses. The frontier now is 2D and 3D models with rotation. The pair instability might be partially suppressed or shifted to higher masses by large stellar rotation, but on the other hand rotationally induced mixing can greatly increase the helium core mass for a given initial mass, thus changing the mass range where different outcomes occur to lower mass.


Recently, we reported in *Nature* [2] a model of the extremely luminous supernova SN 2006gy, which exploded last year in a flash 100 times brighter than typically observed in supernovae. We developed a model that allows us to explain that the brilliant flash of light emanating from the star did not come during a single collapse that resulted in the death of the star — the standard model for typical supernova behavior. Instead, very massive stars with initial masses from 100 to 140 solar masses can produce repeated supernova-like outbursts. Very brilliant bursts of light occur when subsequent shells of material collide with one another at some distance from the star. The kinetic energy of the lost material is converted into light at close to 100 percent efficiency as compared with normal supernovae, where this efficiency is only 1 percent or less.

While previous research has already suggested instability of electron-positron pairs leading to supernova behavior, the *Nature* article [2] is the first time research has suggested that such repeated instability in pulses creates very bright supernovae. With this new model, multiple pulses of varying strengths are possible, giving rise to a wide variety of behaviors in massive stars until they unleash their final knockout punch. Understanding pair-instability supernovae and their observability is also a topic of the newly funded LDRD-ER project, “The First Cosmic Explosions” (PI: Alexander Heger).

Type I X-ray bursts (XRBs) are the thermonuclear explosion of a thin layer of hydrogen and helium on the surface of a neutron star. The immense gravitational acceleration compresses this fuel to the point of ignition, and the fuel is rapidly burned. The typical recurrence time between outbursts is a few hours, so multiple bursts can be observed from a single source. Very detailed timing information along



**Fig. 2. Composition of the ejecta of the 110 solar mass star as it finally dies. The iron core mass is 2.18 solar masses and, its outer edge is collapsing at 1000 kms<sup>-1</sup>. A hot proto-neutron star will now form, which, depending on the maximum physically allowed neutron star mass and accretion over the next few seconds, may become a black hole. A separate calculation of a rotating 95 solar mass star with similar final helium core mass gave an angular momentum for the iron core of  $4.3 \times 10^{48}$  erg•s, implying a neutron star rotation rate of 2 ms. This is enough angular momentum that rotation is likely to play a role in the final death of the star, perhaps producing a millisecond “magnetar” or a “collapsar.”**



**Fig. 3. Surfaces of constant entropy in a 3D simulation of two-component combustion in an X-ray burst. Colors indicate different levels of entropy. Fast combustion occurs when hot carbon gets in contact with a layer of cooler hydrogen above it. Preliminary 1D simulations indicate that the entire hydrogen layer will be engulfed within less than 1 ms. This image shows a simulation of 1 m<sup>3</sup> at the interface between hydrogen and carbon when they first get in contact. The thickness of the entire layer is about 10 m (credit Sanjib Gupta T-16).**

with light curves are provided by the Rossi X-ray Timing Explorer satellite. This frequency of bursts and the availability of good instruments has led to a host of discoveries, including the existence of oscillations during the burst. The likely cause of these oscillations is bright spots on the surface on the rapidly rotating (300-600 Hz) neutron star. These oscillations suggest that the burning is spreading across the neutron star. At LANL XRBs are also studied as part of a large new directed research project on cosmic explosions (PI: Sanjay Reddy, Co-I: Alexander Heger).




**For more information contact Alexander Heger at [aheger@lanl.gov](mailto:aheger@lanl.gov).**

[1] A. Heger and S. E. Woosley, *Astro phys. Jour.*, **567**, 10:532-543 (2002).

[2] S. E. Woosley, S. Blinnikov and A. Heger, *Nature*, **450** (7168), 390-392 (2007).

## Funding Acknowledgments

- Department of Energy, Office of Science, Scientific Discovery through Advanced Computing Program.



# Atomic, Nuclear, and High Energy Physics

This section contains ten articles on fundamental and applied research. These are examples of research in the areas of atomic physics (such as very cold atoms), nuclear physics (such as the calculation of cross-sections for nuclear events), and high energy physics (such as studies of the quark-gluon plasma). The work spans the gamut from purely theoretical to model development to computer simulations.

In atomic physics there are two articles dealing with cold atoms. In one of them, LANL theorists describe how the superfluid energy gaps for lithium atoms can be the largest superfluid gaps for any system. This work is used to interpret experiments carried out at MIT. Another article on cold atoms also relates their properties to those of nuclei. In fact, this relationship arises from universal behavior. A third article in this group deals more with molecular than atomic properties. In it, the authors discuss LEDs made of organic semiconductors, which have applications for efficient lighting.

There are a couple of examples of high-energy physics work in this section. One of the topics is Bose-Einstein condensation, a phenomenon that has been studied mostly in the atomic physics arena. However, in this example we read about Bose-Einstein condensates in a relativistic boson gas at high energies. This work is relevant to experiments at the Relativistic Heavy Ion Collider at Brookhaven National Laboratory. A second article deals with properties of the quark-gluon plasma. The work explains discrepancies between quantum

chromodynamics, the fundamental theory of strong interactions, and recent experimental data.

The most common subject of articles in this section is nuclear physics, an area of prominence for the mission of the Laboratory. One article summarizes work on the calculation of actinide emission cross-sections in nuclear reactions, where a predictive evaluation of pre-equilibrium emission has been achieved. In another article, there is a discussion of the theory of nuclear masses in nuclear reactions. This work has an impact on simulation databases employed across the world. In yet another article is a description of work in simulations of the decay chain of primary nuclear fission fragments. These simulations employ the Monte Carlo method, which was invented at LANL in the 1950s. On a somewhat more theoretical bent, work on the implications of effective field theory on the relativistic dynamics of neutrons and protons is described. Another article deals with potential probes of the fine structure constant, one of the fundamental constants in nature. These measurements involve laser-based investigations of long-lived nuclear isomeric states.



# Bose-Einstein Condensation in the Relativistic Ideal Bose Gas

George A. Baker, Jr., T-11

Currently there is considerable interest in the results being produced by the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory. There, high density quantum chromodynamics (QCD) is very important. In particular the behavior of  $\pi$  mesons and diquarks, which are massive bosons, is relevant. In our just-published paper [1] we study the behavior of the Bose-Einstein condensation for an ideal Bose gas in this regime.

Since its theoretical prediction by Einstein in 1925 based on the 1924 work by Bose on photons, and after many decades languishing as a mere academic exercise in textbooks, Bose-Einstein condensation (BEC) has been observed in the laboratory in laser-cooled, magnetically trapped ultra-cold bosonic atomic clouds. We study instead the conditions in high-density QCD where BEC may occur.

In early papers on the relativistic ideal Boson gas, critical transition temperature  $T_c$ -formulae were derived for both relativistic and ultra-relativistic cases, but the production of anti-bosons was not considered. At high temperature, such as would be expected at RHIC, there will be very considerable production of anti-bosons.

In our work we exhibit, as a function of boson number density, exact BEC transition temperatures for the relativistic ideal boson gas (RIBG) system with and without anti-bosons in 3D. The system with both kinds of bosons always has the higher  $T_c$ , i.e., is the system with the first BEC singularity that appears as it is cooled. This suggests that the Helmholtz free energy might be lower and thus correspond thermodynamically to the *stable* system as opposed to a *metastable* system for the lower- $T_c$  system. It is then calculated and indeed found to be lower, for all densities for the complete problem with both bosons and anti-bosons, when compared with the problem without anti-bosons. This implies that the omission of anti-bosons will not lead to stable states.

The number of bosons  $N$  of mass  $m$  that make up an ideal boson gas in  $d$  dimensions (*without* anti-bosons) is given by

$$N = \sum_{\mathbf{k}} [e^{\beta(|E_{\mathbf{k}}| - \mu)} - 1]^{-1} \quad (1)$$

where  $\beta = 1/k_B T$ ;  $k_B$  is the Boltzmann constant, and  $\mu(T)$  is the boson chemical potential. Here, the total energy of each boson is

$$|E_{\mathbf{k}}| \equiv \sqrt{c^2 \hbar^2 k^2 + m^2 c^4} \quad (2)$$

$$= mc^2 + \hbar^2 k^2 / 2m + O(k^4) \quad \text{if } c\hbar k \ll mc^2 \quad \text{NR} \quad (3)$$

$$= c\hbar k [1 + \frac{1}{2}(mc/\hbar k)^2 + O(k^{-4})] \quad \text{if } c\hbar k \gg mc^2 \quad \text{UR} \quad (4)$$

where  $k$  is the particle wave number,  $m$  refers to the boson rest mass, and  $c$  is the speed of light. The two limits refer to the nonrelativistic (NR) and ultrarelativistic (UR) extremes. For a cubic box of side length  $L$  in the continuous limit the sum in (1) over the  $d$ -dimensional wave vector  $\mathbf{k}$  becomes an integral, namely

$$\sum_{\mathbf{k}} \rightarrow (L/2\pi)^d \int d^d k. \quad (5)$$

At the BEC critical transition temperature  $T_c$ ,  $\mu(T_c) = mc^2$  and the boson number, density can be expressed as

$$n \equiv \frac{N}{L^d} = \frac{1}{(2\pi)^d} \int d^d k \frac{1}{\exp[\beta_c(|E_{\mathbf{k}}| - mc^2)] - 1} \quad (6)$$

where  $\beta_c \equiv 1/k_B T_c$ .

At sufficiently high temperatures such that  $k_B T \gg mc^2$  boson-anti-boson pair-production will occur abundantly. In fact, the total energy  $E_{\mathbf{k}}$  of each particle satisfies  $E_{\mathbf{k}}^2 = c^2 \hbar^2 k^2 + m^2 c^4$  so that

$$E_{\mathbf{k}} = \pm \sqrt{c^2 \hbar^2 k^2 + m^2 c^4} \equiv \pm |E_{\mathbf{k}}| \quad (7)$$

with the  $+$  sign referring to bosons and the  $-$  sign to anti-bosons. The *complete* number equation is now

$$\begin{aligned} N - \bar{N} &\equiv \sum_{\mathbf{k}} (n_{\mathbf{k}} - \bar{n}_{\mathbf{k}}) \\ &= \sum_{\mathbf{k}} \left[ \frac{1}{\exp[\beta(|E_{\mathbf{k}}| - \mu)] - 1} - \frac{1}{\exp[\beta(|E_{\mathbf{k}}| + \mu)] - 1} \right] \end{aligned} \quad (8)$$



where  $n_{\mathbf{k}}$ , ( $\bar{n}_{\mathbf{k}}$ ) is the average number of bosons (anti-bosons) in the state of energy  $\pm |E_{\mathbf{k}}|$ , respectively, at a given temperature  $T$  and  $N$  ( $\bar{N}$ ) is their respective total number at that temperature. Since  $n_{\mathbf{k}}, \bar{n}_{\mathbf{k}} > 0$  for all  $\mathbf{k}$  and  $E_0 = mc^2$ , the chemical potential must be bounded by

$$-mc^2 \leq \mu \leq mc^2. \quad (9)$$

Instead of  $N$  constant we must now impose the constancy of  $N - \bar{N}$  to extract the correct BEC critical temperature, which we call  $T_c^{B\bar{B}}$  as it involves both bosons ( $B$ ) and anti-bosons ( $\bar{B}$ ). At  $T = T_c^{B\bar{B}}$  one has  $|\mu(T_c^{B\bar{B}})| = mc^2$ . Figure 1 shows our results.


**For more information contact George Baker at [gjb@lanl.gov](mailto:gjb@lanl.gov).**

#### Reference

[1] M. Grether, M. de Llano, and G.A. Baker, Jr., *Phys. Rev. Lett.* **90**, 200406 (2007), and the references therein.

#### Funding Acknowledgments

- Department of Energy, Office of Science



# Microscopic Calculation of the Pre-equilibrium Emission for Neutron Scattering on Actinides

Marc Dupuis, Ludovic Bonneau, Toshihiko Kawano, T-16

When a nucleon collides with a target nucleus, several reactions may occur: elastic and inelastic scatterings, charge exchange, multiple particle emissions, etc. To explain the particle emissions, three nuclear reaction processes are involved: the direct process, the pre-equilibrium or pre-compound emission, and the formation and evaporation of a compound nucleus. Moreover, for heavy elements like actinides ( $^{238}\text{U}$ ,  $^{232}\text{Th}$ , etc.), a fission process takes place and more particles are emitted. Here we focus on the neutron-induced reactions on  $^{238}\text{U}$  in the energy range 10–20 MeV. To predict the energy spectrum and the angular distribution of the emitted neutrons, one has to use several nuclear reaction models to take into account the different reaction processes. However, these models rely on adjustable parameters, which are usually tuned to reproduce the experimental data. Consequently, the evaluation of the emission spectrum remains ambiguous, and no reliable extrapolation can be made to domains when there is no experimental data. The example of neutron scattering on actinides is very important—the experimental data are scarce, but accurate evaluations are needed for applications like neutron moderation in nuclear reactors.

Hence, more predictive evaluations are really needed. For this purpose, we have achieved a calculation of the pre-equilibrium emission that does not require any adjustment to the experimental data. The cross-section for neutron emission is computed from all possible transitions between the target ground state and excited states, which follow the collision between the target and the projectile. These transitions are calculated within a quantum mechanical framework—wave functions of the target states are obtained with a deformed Hartree-Fock calculation, and we use a nuclear two-body interaction obtained from microscopic calculation to generate the transitions. We present in Fig. 1 the energy distributions obtained for the scattering of 14 MeV neutrons on the  $^{238}\text{U}$  target. For the same reaction, we depict in Fig. 2 the angular distribution of 7.5 MeV outgoing neutrons. On these figures, we have added the contribution coming from the evaporation and fission processes. Our pre-equilibrium calculations reproduce the energy and angular distribution very well without any fitting to the experimental data—this is good progress in the area of pre-compound reaction modeling.




Fig. 1. Angle integrated neutron emission spectrum for 14 MeV neutron scattering on  $^{238}\text{U}$ .

Microscopic calculations will also be needed to describe the other processes to get a precise prediction of the whole energy spectrum. Indeed, the models of compound nucleus and fission use many phenomenological parameters. Better calculations of direct processes, which contribute to the high energy part of the neutron spectrum, are also needed since they are performed with phenomenological interactions and a crude description of the target states. More predictive calculations of the fission and the direct processes could be performed with an accurate description of the target ground and excited states. A precise structure description of heavy elements is still very difficult, but the new computer capacities allow us to consider such calculations.




Fig. 2. Angular distribution of emitted neutron at 7.5 MeV for 14 MeV neutron scattering on  $^{238}\text{U}$ .

For further information contact Marc Dupuis at [marcd@lanl.gov](mailto:marcd@lanl.gov).

#### Funding Acknowledgments

- Department of Energy, Office of Science
- Department of Energy, Office of Science, Nuclear Physics Program
- Department of Energy, Office of Science, Scientific Discovery through Advanced Computing Program
- United States Nuclear Data Program

# A Relativistic Symmetry in Nuclei and Effective Field Theory

Joseph N. Ginocchio, T-16

**T**he dynamics of neutrons and protons in nuclei have been successfully treated nonrelativistically. Therefore it came as a surprise that the (almost) equality of certain pairs of energy levels of quantum states (quasi-degeneracy) in heavy nuclei can be explained by a relativistic symmetry [1]. This symmetry is called pseudospin symmetry because it has the same mathematical properties as spin but different physical properties. For example, the quasi-degenerate quantum states have different radial quantum numbers and different orbital angular momenta, features not possible with spin, which made the reason for their quasi-degeneracy difficult to penetrate. Pseudospin entangles both spin and momentum of the nucleons in the nucleus in such a way to make such transitions possible.

Since pseudospin symmetry seems to be more valid than spin in nuclei, the question arises if there is a fundamental rationale for this symmetry. Recently effective field theories have been developed that derive effective nucleon-nucleon interactions. These effective nucleon-nucleon interactions involve the spin. We may ask why don't these expansions naturally involve the pseudospin as well? We have investigated this issue and discovered some interesting answers.

First we found that the spin-spin interaction between nucleons is equivalent to the pseudospin-pseudospin interaction between nucleons. Second, we showed that the tensor interaction is a spin-pseudospin interaction. This is an interesting insight and implies that the tensor interaction violates both spin and pseudospin equally. These results demonstrate that effective field theories put spin and pseudospin on an equal footing but, by the same token, there is no need to introduce pseudospin.

However, the two body spin-orbit and the two body pseudospin-pseudo-orbit interactions are not equivalent and imply different physics. The two body pseudospin-pseudo-orbit interaction can be written in terms of a linear combination of the two body spin-orbit interaction and the tensor interaction. Although the tensor interaction conserves neither spin or pseudospin, the two body spin-orbit interactions conserves spin but not pseudospin and, vice-versa, the two body pseudospin-pseudo-orbit interaction conserves pseudospin but not spin.

This suggests that, instead of writing the effective nucleon-nucleon interaction in terms of the tensor interaction, a more revealing exposition would be to write the effective nucleon-nucleon interaction as a linear combination of the two body spin-orbit interaction and two body pseudospin-pseudo-orbit interaction. If the strength of the two body pseudospin-pseudo-orbit is much larger than the strength of the two body spin-orbit interaction, then the nucleon-nucleon interaction would have an approximate dynamical pseudospin symmetry and would explain the observance of approximate dynamical pseudospin symmetry in nuclei. This program is being carried out.

**For more information contact Joseph Ginocchio at [gino@lanl.gov](mailto:gino@lanl.gov).**

[1] Joseph N. Ginocchio, *Physics Reports* **414**, 165 (2005).

## Funding Acknowledgments

- Department of Energy



## Sensitivity of Nuclear Transition Frequencies to Temporal Variations in the Fine Structure Constant ( $\alpha$ -dot)

Anna C. Hayes, James L. Friar, T-16

There exist in nature a few long-lived nuclear states (isomers) with very low (eV) excitation energies. The best-known example is the 7.6 eV isomer of  $^{229}\text{Th}$ , which is estimated to have a lifetime of about 5 hours for decay to the ground state. The combination of the low energy and narrow width for the isomer makes this transition a candidate for laser-based investigations and possibly a sensitive probe for any temporal variation in the fine-structure constant  $\alpha = e^2/\hbar c$ .

Some evidence for a temporal variation of the fine structure constant has been reported from an analysis of quasar absorption spectra that suggests  $\dot{\alpha}/\alpha \sim 7.2 \times 10^{-16}/\text{year}$ . Laboratory searches for  $\dot{\alpha}/\alpha$  have concentrated on measurements of temporal changes in transition frequencies between nearly degenerate atomic levels. More recently it has been suggested that the nearly degenerate nuclear levels in  $^{229}\text{Th}$  may provide several orders of magnitude more sensitivity to  $\dot{\alpha}/\alpha$  than do atomic levels (Fig. 1). In the present research we examined the underlying physics that determines the sensitivity of nuclear transition frequencies to  $\dot{\alpha}/\alpha$ .

The energy of a transition ( $\omega$ ) between the ground state and an excited state of any nucleus is determined by the difference of matrix elements of the nuclear Hamiltonian, and the sensitivity to  $\alpha$  enters through the Coulomb interaction between nucleons in the nucleus. Invoking an extension of the Feynman-Hellmann theorem, we found the simple result that a temporal variation of the transition frequency  $\omega^*$  depends on  $\dot{\alpha}/\alpha$  as:


$$\omega^* = (\langle f | V_C | f \rangle - \langle g.s. | V_C | g.s. \rangle) \dot{\alpha}/\alpha$$

where  $\langle f | V_C | f \rangle$  and  $\langle g.s. | V_C | g.s. \rangle$  are the isomeric-state and ground-state Coulomb energies. Thus the variation in the transition frequency due to the fine-structure constant is driven solely by the Coulomb-energy difference of the two states.

The low-lying states in  $^{229}\text{Th}$  are normally described in terms of the nuclear Nilsson model. In this picture the states are modeled as an active neutron in a deformed potential well. This model predicts that the Coulomb energies of the doublet are the same, and that they arise from the same deformed nuclear core. Thus  $(\langle f | V_C | f \rangle - \langle g.s. | V_C | g.s. \rangle) = 0$  and there is no sensitivity to  $\dot{\alpha}/\alpha$ . We also examined a more sophisticated model for  $^{229}\text{Th}$  that represents the first-order correction to the Nilsson model, namely the inclusion of the pairing interaction between neutrons. However, in this model the sensitivity to  $\dot{\alpha}/\alpha$  remained zero because the Coulomb deformation of the two states was unchanged.

We concluded that a large enhancement in the temporal change of the nuclear transition frequency due to a finite  $\dot{\alpha}/\alpha$  is difficult to realize in  $^{229}\text{Th}$ . The main issue is that the magnitude of the enhancement is determined by the difference of the Coulomb energies of the two states, and that the latter is constrained to be small because the two states differ mainly in their neutron (but not their proton) structure.





For further information contact A. C. Hayes at [anna\\_hayes@lanl.gov](mailto:anna_hayes@lanl.gov)

**Funding Acknowledgments**

- Department of Energy, Office of Science
- Los Alamos National Laboratory Directed Research and Development Program

# Nuclear Masses for Simulation Databases

Peter Möller, Arnold J. Sierk, T-16

**N**uclear masses determine the energy balance in nuclear reactions. For example, in radioactive decay such as  $\alpha$  and  $\beta$  decay, in element synthesis reactions such as neutron capture in stars, and in heavy-ion reactions in laboratories. In a nuclear reactor, power is generated when uranium fissions into two fragments. Many different mass and charge splits may occur, leading to hundreds of different isotopes of about 40 different elements. To understand the properties of a reactor, a model of these decays is needed. The energy released in, and the half-life of a decay from a parent nucleus to one or more daughters are governed by the relative masses of the parent and daughter(s). Some of the masses involved have been measured, but many have not, due to their short lifetimes. Masses of such short-lived or hard-to-produce isotopes have to be obtained from a model. The most used nuclear mass model in the world was developed at Los Alamos. It involves calculating the shape and energy of the most stable nuclear configuration for each isotope (its ground state). The model takes microscopic effects into account by solving the quantum-mechanical Schrödinger equation on a cluster of computers, leading to masses for any proton number  $Z$  and neutron number  $N$ . A first model was published in 1981 and an improved one in 1995 [1]. A measure of the importance of this work is that it is the most downloaded paper from this journal 12 years after publication, and that the model tables are the preferred or only database for theoretical nuclear masses at, for example, the National Nuclear Data Center at Brookhaven National Laboratory and the International Atomic Energy Agency. Our paper has also been one of the three or so most cited LANL papers since its publication in 1995.

A goal for a theory of nuclear masses is that it be able to accurately predict masses of nuclei for which no measured values are available. Our finite range droplet model (FRDM) (1992), published in 1995, was adjusted to a 1989 database of nuclear masses. In 1997 we compared our published masses to 217 experimental masses measured after 1989 and found that the average error was unchanged (Fig. 1). The most recent experimental database has 529 “new” masses measured since 1989. The error of our masses with respect to the new data [2] is 0.46 MeV, much smaller than with respect to the 1989 data, (Fig. 2). An interesting observation is that some points on the proton-rich side of  $\beta$  stability indicating large deviations




Fig. 1. Our mass model FRDM (1992) was adjusted to a 1989 database of experimental masses. We here compare the model predictions to masses measured after 1989 up to 1993. No increase in the model error is observed. The light gray area is the model one-sigma error for nuclear masses to which the model was adjusted.

between theory and experiment in Fig. 1 are not present in Fig. 2. Since the theory is unchanged, the difference means that several experimental masses were either removed from the evaluated database, or changed by more than the error bars!

Another way to study the model reliability is to plot the difference between calculated and experimental masses for each proton number  $Z$  and neutron number  $N$ . In Fig. 3, we show the differences with respect to the 1989 data to which the model was adjusted. In Fig. 4, we show the differences with respect to all the 2003 masses. Significantly, the previously unknown nuclei along the upper and lower edges of the colored region in Fig. 4 exhibit very small differences between the model and recently measured masses.

Because the computer resources available to us are 100,000 times more powerful than when we started our work on the FRDM (1992) in 1987, we have undertaken a more accurate and sophisticated determination of ground-state shapes. This work is almost completed, leading to a mass table with an average error of about 0.585 MeV (compared to the previous average error of 0.669 MeV).




Fig. 2. Since 1993 an additional 300 masses have been measured so we can now compare our model predictions to more than 500 masses that were not known to us when the model calculations were published. The predictive power is excellent and the error for the predicted masses is considerably smaller than in the region where the model was adjusted. Moreover, masses on the proton-rich side, which in Fig. 1 exhibit some of the largest deviations from our calculated masses, are no longer present in this figure. This means they have been found in error and removed from the evaluated database, or the measured values have been revised by more than the error bars.

For further information contact Peter Möller at [moller@lanl.gov](mailto:moller@lanl.gov).

[1] Atomic Data and Nuclear Data Tables **59**, 185 (1995).

[2] E. Pinheiro, W.D.Weber, and L.A. Barroso, In *Proceedings of the 1995 International Conference on Parallel Processing (ICPP)* (1995).

#### Funding Acknowledgments

- Department of Energy, National Nuclear Security Administration, Advanced Simulation and Computing Program
- Los Alamos National Laboratory Directed Research and Development Program




Fig. 3. Calculated masses compared to those experimental masses to which the model parameters were adjusted.




Fig. 4. Calculated masses compared to all experimental masses in the most recent evaluation. More than 500 masses were not known when the model was presented. The model gives excellent predictions of masses for these nuclei.

# Prompt Neutrons and Fission Modes

Patrick Talou, Toshihiko Kawano, T-16

In our understanding of the nuclear fission process, what happens near the scission point where the two nascent fragments separate remains unclear. Observations of the final fission products in terms of charge, mass, and kinetic energy distributions provide indirect clues about the process. Another indirect but closer probe of the physics near scission is to look at the characteristics of the neutrons emitted by the fission fragments right after their separation. These so-called prompt neutrons are emitted from the excited primary fragments until they settle in a stable state of a nucleus (before eventually undergoing a subsequent beta-decay).

We have studied the consecutive emissions of prompt neutrons and gamma-rays by simulating the decay chain of the primary fission fragments. Each stage (i.e., each neutron emission) of this decaying sequence is treated exactly, assuming an evaporation process from a temperature-dependent Weisskopf spectrum. This type of simulation provides a large set of data, and can help reveal interesting correlations among the emitted neutrons (e.g., in energy, angle, multiplicity).

As an example, Monte Carlo simulations were performed for the low-energy neutron-induced fission of  $^{235}\text{U}$ , for incident energies varying from 0.5 to 6.0 MeV. As an input to our simulations, the recently measured fission fragment yields  $Y(A, TKE)$  from Hambsch were used [1]. Here,  $A$  represents the mass of the fission fragment, and  $TKE$  is the total kinetic energy of the two complementary fragments. These yields were fitted using fission modes as defined in the Brosa model [2]. Roughly speaking, these fission modes originate from different paths leading to fission taken by the nuclear system along a multi-dimensional potential energy surface.

The experimental yields  $Y(A, TKE)$  at 0.5 MeV incident neutron energy are depicted in Fig. 1, along with their decomposition in fission modes. As an example, the dominant S2 mode ( $\sim 75\%$  of the total) corresponds to an asymmetric fission

(a light fragment and a heavy fragment are created) and to medium-to-high  $TKE$  values, corresponding to relatively compact shapes at scission. On the contrary, the SL mode corresponds to very elongated shapes with very low  $TKE$  values.

Monte Carlo simulations of the decay of the excited primary fragments were performed by sampling the  $Y(A, TKE)$  distribution for each fission mode. Figure 2 shows the average distributions of the neutron multiplicity  $\langle \nu \rangle$  as a function of mass  $A$  and total kinetic energy  $TKE$ , for each fission mode. By projecting on the mass axis, we get the result shown in Fig. 3, and compared to available experimental data.

Only this level of detail in the calculations can help us better understand the physics near the scission point, and constrain theoretical models and physics assumptions. One particularly important question being addressed is: how is the total excitation energy available at scission shared between the two nascent fragments? In the particular situation shown here, the light fragment takes a larger share of this energy than its heavy counterpart, principally due to the difference in deformation energies of the fragments at scission.

**For further information contact Patrick Talou at [talou@lanl.gov](mailto:talou@lanl.gov).**

- [1] F.J. Hambsch, IRMM Geel, Belgium, private communication (2007).
- [2] U. Brosa, S. Grossmann, and A. Müller, *Phys. Rep.*, **197**, 167 (1990).
- [3] P. Talou, to appear in *Conf. Proc. of the Int. Conf. on Nuclear Data for Science & Technology ND2007*, Nice, France, April 22-27 (2007).

## Funding Acknowledgments

- Department of Energy, National Nuclear Security Administration, Advanced Simulation and Computing Program




Fig. 1. Experimental yields  $Y(A, TKE)$  of  $n(0.5 \text{ MeV}) + {}^{235}\text{U}$  and their decomposition in Brosa fission modes S1, S2, and SL. The top right plot labeled SUM corresponds to the full experimental distribution, i.e., the weighted sum of all three modes with the S2 mode representing about 75% of the total. TKE is the total kinetic energy of the two fragments.




Fig. 2. Average prompt neutron multiplicity  $\langle \nu \rangle$  (Fragment mass, TKE) for each fission mode.




Fig. 3. Projection of results shown in Fig. 2 on the fission fragment mass axis and compared with existing experimental data.



# Novel Heavy Flavor Suppression Mechanisms in the Quark-gluon Plasma

Ivan Vitev, T-16

Quantum chromodynamics, the fundamental theory of strong interactions, predicts that at exceedingly high temperatures and densities ordinary nuclear matter should melt down in a plasma-like state of quarks and gluons. One of the important experimental signatures of the quark-gluon plasma (QGP) creation in heavy ion collisions is the detailed suppression pattern of high transverse momentum ( $p_T$ ) hadrons, dubbed jet quenching [1]. Jet quenching for light mesons, such as  $\pi$ , K, and  $\eta$ , at the Relativistic Heavy Ion Collider (RHIC) is well explained by radiative energy loss calculations. One such example of neutral pion suppression in Au+Au and Cu+Cu reactions is shown in Fig.1. Our theoretical calculations of the QGP-induced suppression  $R_{AA}$  versus  $p_T$  and the colliding system size show good agreement with the PHENIX experimental data.

In contrast, recent data on the suppression of single non-photonic electrons, coming from the decay of heavy charm (D) and beauty (B) mesons, cannot be reproduced by medium-induced gluon bremsstrahlung phenomenology in conjunction with a physically reasonable set of QGP temperatures and densities. Recently, we identified a possible reason for this striking discrepancy [2]. The model for predicting the quenched hadron spectra in A+A collisions, where final-state-in-medium effects are important, has been very simple: it is assumed that the hard jet hadronizes in vacuum, having fully traversed the region of dense nuclear matter,  $L_{QGP}^T \sim 6$  fm. It is easily demonstrated that this assumption is not valid for the heavy charm and beauty quarks and mesons. In fact, for  $p_T = 10$  GeV they tend to form at times  $t = 0.4$  fm and  $t = 1.5$  fm, respectively, well within the hot and dense QGP phase. Consequently, the natural mechanism of heavy flavor suppression is a competition between D- and B-meson dissociation and c- and b-quark fragmentation, both of which emulate energy loss. Our predictions for the suppression  $R_{AA}(p_T)$  of  $e^+e^-$  coming from the semi-leptonic heavy flavor decays are compared in Fig. 2 with the STAR and PHENIX data [3]. Thus, our approach presents an elegant resolution to the “heavy flavor puzzle” at RHIC.

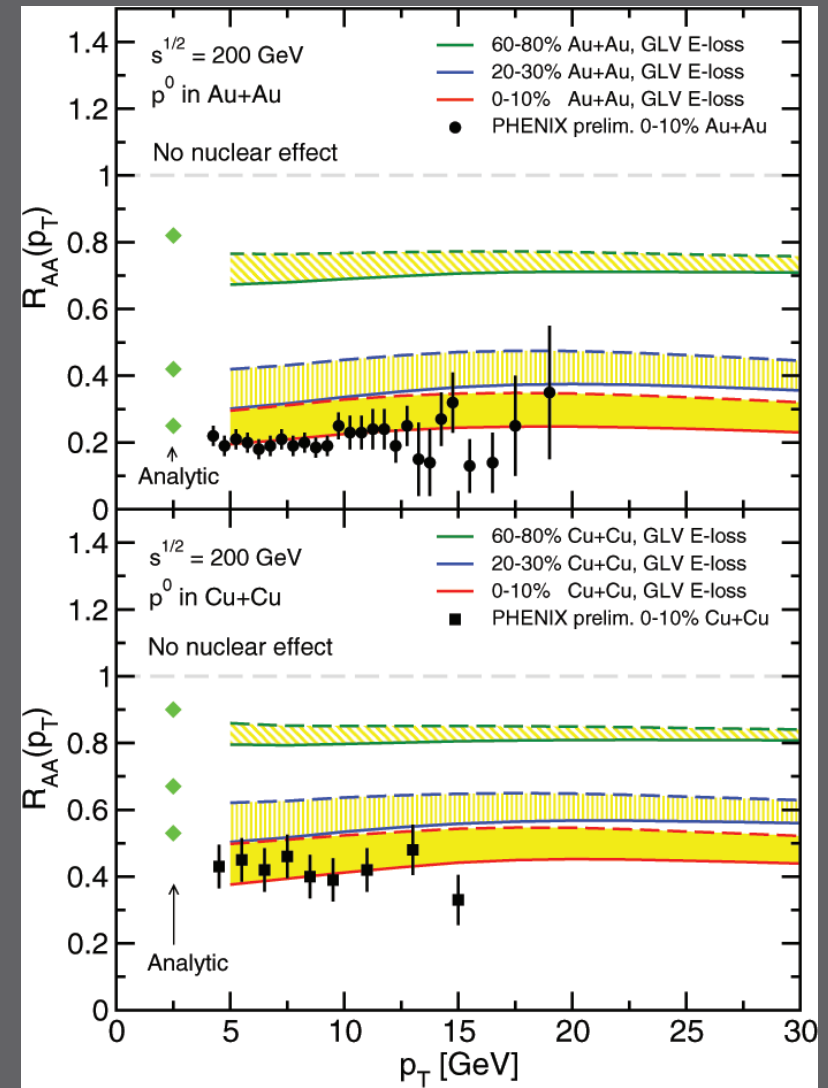


Fig. 1. Top panel: neutral pion quenching predictions versus  $p_T$  and centrality in Au+Au collisions at RHIC compared with the PHENIX central suppression data at C.M. energy of 200 GeV. Bottom panel: similar comparison is done in Cu+Cu reactions.

A natural *unique* consequence [4] of this theory is that B-mesons are attenuated as much as D-mesons at transverse momenta as low as  $p_T \sim 10$  GeV for the conditions prevalent in central Au+Au collisions at C.M. energy of 200 GeV (Fig. 3). Our model will be put to a critical test by experimental advances aiming at direct and separate measurements of D- and B-meson  $R_{AA}(p_T)$  [5].

For further information contact Ivan Vitev at [ivitev@lanl.gov](mailto:ivitev@lanl.gov).

- [1] M. Gyulassy et al., *Quark-Gluon Plasma III*, 123, Eds. R.C. Hwa and X.-N. Wang, World Scientific (2004).
- [2] I. Vitev, *Phys. Lett. B*, **639**, 38 (2006).
- [3] A. Adil and I. Vitev, *Phys. Lett. B*, **649**, 149 (2007).
- [4] I. Vitev, A. Adil, and H. van Hees, *G. Phys. G.*, **34**, S769 (2007).
- [5] P. McGaughey et al., "Heavy quarks as a probe of a new state of matter", LDRD 20060049DR.

#### Funding Acknowledgments

- Department of Energy, Office of Science
- J. Robert Oppenheimer Fellowship
- Los Alamos National Laboratory Directed Research and Development Program




Fig. 2. Suppression of inclusive non-photonic electrons from D- and B-meson spectra softened by collisional dissociation in central Au+Au collisions. Data on non-photonic electron quenching is from PHENIX and STAR.




Fig. 3. Individual suppression of D- and B-meson production via collisional dissociation in the QGP. Top panels show numerical results for  $R_{AA}(p_T)$  in central Au+Au and Cu+Cu collisions at RHIC. Bottom panel shows predictions for central Pb+Pb collisions at the LHC.

# Organic Semiconductor Light-emitting Diodes for Solid State Lighting

Ping Yang, Enrique R. Batista, Sergei Tretiak, Richard L. Martin, T-12; Avadh Saxena, Darryl L. Smith, T-11

The object of solid state lighting is to replace conventional incandescent and fluorescent light bulbs with highly efficient solid state light-emitting diodes (LEDs). Lighting is an extremely large business that consumes about one quarter of the electricity produced annually in the United States. Over  $10^{18}$  J (about 300 Megatons of TNT) of energy are expended for lighting annually. Efficient semiconductor LEDs have the potential to replace both fluorescent and incandescent lights as the primary lighting source. They are potentially more efficient and more environmentally benign (no Hg) than these current lighting sources. In addition, semiconductor LEDs have unique properties that make them attractive for lighting applications—it is possible to electrically control the spectral properties of the light emitted, and they can be arranged over large areas in various shapes for aesthetic purposes.

There are two basic kinds of semiconductor LEDs being explored for solid state lighting: inorganic semiconductor LEDs, and organic semiconductor LEDs. Inorganic semiconductor LEDs use crystalline materials like GaN and organic semiconductor LEDs use disordered materials based on conjugated hydrocarbons (i.e., materials with double carbon-carbon bonds), like the polymer poly-phenylene vinylene (PPV). Inorganic semiconductor LEDs make bright point light sources. Figure 1 shows a set of inorganic semiconductor white LEDs used for solid state lighting. Organic semiconductor LEDs make large area light sources. Figure 2 shows a white organic LED illuminating a color chart. At present inorganic LEDs are somewhat further developed than organic LEDs for solid state lighting applications.

At Los Alamos we are involved in the investigation of organic semiconductors for a variety of applications including organic LEDs for solid state lighting. The basic device structure of an organic LED is shown in Fig. 3. It consists of a thin (typically about 100 nm) layer of organic semiconductor sandwiched between two conducting contacts. The organic semiconductor is undoped and highly resistive. Electrons are injected into the organic semiconductor from one of the conducting contacts and holes are injected into the organic semiconductor from



Fig. 1. Inorganic semiconductor white LEDs for solid state lighting.




Fig. 2. Organic semiconductor white LEDs for solid state lighting.




Fig. 3. Basic device structure of an organic LED.




Fig. 4. Top: molecular structure of a five-ring oligomer of PPV. Middle, left: crystal structure. Middle, right: disordered structure of PPV oligomers. Bottom: calculated density of states for: A) crystal, B) disordered cluster, C) isolated ordered oligomer, and D) ensemble of isolated disordered oligomers. From [1].

the other conducting contact. The electrons and holes recombine in the organic semiconductor, emitting light whose color is determined by the energy gap of the organic semiconductor. The conducting contacts have different work functions—a high work function contact to inject holes, and a low work function contact to inject electrons. One of the conducting contacts must be transparent to let the light out of the structure, and indium tin oxide is often used for this purpose.

In our experiment's collaboration with colleagues from MPA division (Ian Campbell and Brian Crone), we are studying the basic properties of organic semiconductor devices such as LEDs. We investigated the role of intra-molecular conformational disorder and inter-molecular electronic interactions on the electronic structure of disorder clusters of (PPV) oligomers. Classical molecular dynamics was used to determine probable molecular geometries, and first-principle density functional theory (DFT) calculations were used to determine electronic structure. Intra-molecular and inter-molecular effects were disentangled by contrasting results for densely packed oligomer clusters with those for ensembles of isolated oligomers with the same intra-molecular geometries. We found that electron trap states are induced primarily by intra-molecular configuration disorder, while the hole trap states are generated primarily from inter-molecular electronic interactions. Figure 4 shows an example of these calculations, which was recently published [1].

Organic semiconductors have proven to be technologically valuable for a number of applications. Although the device details depend on the envisioned application, the basic physical processes involved do not. Therefore, a fundamental understanding of the physical processes that take place in organic semiconductors is important for a wide variety of applications. The most developed technological application of organic semiconductors at this time is displays. Figure 5 shows a flat panel TV made from organic LEDs that Sony has just introduced to the commercial market.

**For further information contact**  
Darryl L. Smith at [dsmith@lanl.gov](mailto:dsmith@lanl.gov).

[1] P. Yang et al., *Phys. Rev. B*, **76**, 241201R (2007).

#### Funding Acknowledgments

- Department of Energy, Office of Science, Office of Basic Energy Sciences



Fig. 5. Flat panel TV made from organic LEDs. Manufactured by Sony.



# The Largest Superfluid Gaps in Nature and in the Laboratory

Sanjay Reddy, Alexandros Gezerlis, Joseph Carlson, T-16

Theorists in Los Alamos predicted that the largest pairing gap would be found in cold atom experiments with trapped  $^6\text{Li}$  atoms [1]. In cold atom experiments it is possible to adjust the strength of the interaction, and in particular to tune the interaction so that two isolated particles would be on the threshold of being bound—the so-called unitary regime where the scattering-length between particles becomes infinite. Consequently, in the unitary regime the only physical scale in a many-particle system is the Fermi momentum or the inverse inter-particle distance. All physical quantities, including the ground-state energy and the pairing gap, must be universal constants times the Fermi energy of the non-interacting system. Previously, LANL theorists had predicted the ground-state energy, and this result has been confirmed by experimentalists at MIT, Rice, and Duke universities.

LANL researchers had also predicted the pairing gap in cold atoms at unitarity to be half the Fermi energy. This is in stark contrast to traditional superfluids and superconductors, where the pairing gap is thousands of times smaller than the Fermi energy. However, a direct and unambiguous method to measure the superfluid-gap in cold-atom experiments did not exist. To address this problem, Carlson and Reddy analyzed the polarization measurements performed by the MIT group [2] and showed that it can provide very stringent lower and upper bounds on the value of the pairing gap.

By trapping and cooling  $^6\text{Li}$  atoms in two different hyperfine states (labeled spin-up and spin-down) the MIT team was able to observe phase-separation between a superfluid inner-core region with nearly equal spin-up and spin-down particles, and a highly polarized normal state in the outer shell with different spin-up and spin-down densities. The two regions were separated by a first-order transition. Amazingly, in the vicinity of the interface the superfluid phase carried a sizeable polarization at low but finite temperatures (Fig. 1). LANL theorists have analyzed these experiments and found this polarization near the interface is in good agreement with quantum Monte Carlo predictions for the equation of state and the pairing gap. Figure 1 shows the comparison of theoretical results with experimental




Fig. 1. Theoretical predictions for the pairing gap for neutron matter and cold atoms as a function of density or coupling strength.

measurement of the polarization density as a function of the radius of the atomic trap. The polarization of the superfluid in the vicinity of the transition located at  $r \sim 0.4$  requires a superfluid gap in the range of 0.4-0.5 times the Fermi energy. This is dramatically larger than any other fermionic system measured in the laboratory.

Going beyond the laboratory, the largest pairing gap in nature can be found in the exotic astrophysical environments of neutron stars. The interaction between two neutrons is nearly strong enough to bind them, and is therefore quite similar to cold atoms in the unitary regime. In ordinary nuclei, the range of the interaction is comparable to the inter-particle spacing, and the pairing gap is 'only' about 10 percent of the Fermi energy. LANL researchers have examined the equation of state and the pairing gap in low-density neutron matter, as occurs in the crusts of neutron stars. The equation of state is found to be exceedingly close to that of cold atoms. The pairing gap is shown in Fig. 2. It is large, but reduced somewhat compared with cold atom gas because of the finite range of the interaction. LANL






Fig. 2. Polarization versus radius observed in the cold-atom experiment. Comparison between LANL theory and the MIT experiment yields stringent lower and upper bounds on the gap.

researchers find a maximum gap of around 30 percent of the Fermi energy for neutron stars. The gap is also very large compared with the typical temperature inside neutron stars, rising to a maximum of 2 MeV. This large gap can be a very important in astrophysics, affecting the cooling rates of neutron stars. Similar large gaps may also occur in the inner cores of neutron stars where deconfined quarks can form Cooper pairs. The excellent agreement between theory and experiment in the cold-atom context lends credence to extrapolating theory to strongly interacting Fermi systems in neutron stars.

These strongly interacting Fermi systems also provide a rich opportunity to search for exotic states of matter including so-called gapless superfluids, where the groundstate is simultaneously polarized and superfluid. Indeed, LANL researchers have predicted this state will occur at even stronger interactions, where two particles in a vacuum would be slightly bound. Experimentalists are presently searching for this phase of matter.

For further information contact Sanjay Reddy at [reddy@lanl.gov](mailto:reddy@lanl.gov).

- [1] J. Carlson and S. Reddy, *Phys. Rev. Lett.*, **96**, 060401 (2005).
- [2] Y. Shin et al., arXiv:0709.3027v2 (2007).
- [3] J. Carlson and S. Reddy, arXiv:0711.0414v2, accepted by *Phys. Rev. Lett.* (2008).
- [4] A. Gezerlis and J. Carlson, *Phys. Rev. C* **77**, 032801 (R) (2008).

#### Funding Acknowledgments

- Department of Energy, Office of Science, Scientific Discovery through Advanced Computing Program
- Los Alamos National Laboratory Directed Research and Development Program

## From Nuclei to Cold Atoms and Back

Ionel Stetcu, T-16; Bruce R. Barrett, Ubirajara van Kolck, Univ. of Arizona;  
James P. Vary, Iowa State Univ.

**S**ystems with large two-body scattering lengths are of particular interest because they exhibit universal behavior. While such systems have become popular in atomic physics only in the last decade, they have been investigated in nuclear physics since its beginning. Hence, it is not surprising that techniques developed for solving the nuclear few-body problem are immediately applicable to such systems in atomic physics.

The no-core shell model (NCSM) is a powerful many-body technique that has been successfully employed to determine the properties of few-body systems with no adjustable parameters. In this approach, a numerical solution to the many-body non-relativistic Schrödinger equation is obtained by exact diagonalization in a restricted space. Correlations left out by the space restriction are traditionally taken into account by deriving, via a unitary transformation, an effective interaction from two- and three-body forces adjusted to experimental two- and three-nucleon data.

In order to mitigate some shortcomings in the traditional NCSM and, at the same time, provide a transparent connection with the underlying strong interactions at the quark level, we have proposed a novel approach [1] in which the interaction in the restricted NCSM model space is constructed following the general principles of effective field theories (EFTs). The basic idea underlying EFTs is that the restriction of a theory to a model space generates all interactions allowed by the theory's symmetries. Since particle momenta are limited within the restricted space, one can treat short-distance interactions in a derivative expansion, similar to the multipole expansion in classical electrodynamics. The coefficients of this expansion carry information about the details of the short-range dynamics. We have successfully applied this hybrid NCSM and EFT method in the four- and six-body systems, obtaining reasonable agreement with experiment.




Fig. 1. Lowest energy levels for  $j^\pi = 1/2^-$  (circles) and  $j^\pi = 1/2^+$  (squares) states of the three-fermion system in a harmonic trap as function of the ratio  $b/a_2$ . Around  $b/a_2 \sim 1.5$ , we observe an inversion of the parity of the system's ground state. The dashed line is a fit of the positive-parity points by a quadratic form.

The same general principles to construct effective interactions in restricted model spaces can be applied to other systems. Thus, we have applied similar methods for the description of systems of three and four spin  $1/2$  fermions in harmonic traps [2]. This allows us to investigate the properties of few-atoms from the weakly interacting to the strongly interacting limits. Note that, although for low-momenta the effective Hamiltonian for cold atoms and nuclear systems are formally similar, the coefficients involved carry different information. In the three-fermion system, we found that the ground state changes parity when the interaction increases beyond the unitary limit. In Fig. 1, we show the lowest positive and negative  $1/2$  states of the three-fermion system, as a function of the ratio between the trap length  $b$ , and the two-body scattering length  $a_2$ . In the  $b/a_2 \rightarrow \infty$  limit, the positive-parity state likely represents the untrapped  $S$ -wave particle-dimer scattering state that one expects to dominate sufficiently close to threshold.

We further consider the application of the same principles to the construction of the effective interaction in finite NCSM spaces for the description of the nuclear many-body problem. Thus, the formal similarity between the physically trapped fermions and interacting nucleons trapped by a center-of-mass harmonic term allows us to use similar means to obtain the effective interaction more efficiently for few-nucleon calculations.


**For further information contact Ionel Stetcu at [istet@lanl.gov](mailto:istet@lanl.gov).**

[1] I. Stetcu, B.R. Barrett, and U. van Kolck, *Phys. Lett. B*, **653**, 358 (2007).

[2] I. Stetcu, et al., *Phys. Rev. A* (in press), arXiv:0705.4335 (2007).

### **Funding Acknowledgments**

- Department of Energy
- National Science Foundation





## Chemistry and Bioscience

LANL has a long history of internationally known work in the areas of chemistry and bioscience. Much of the research presented here has been funded by “work for others” agencies like the National Institutes of Health and the Office of Science at the Department of Energy.

The chemistry series of articles in this section we start with a discussion of the electronic transport in disordered organic polymers. These are compounds that have applications in electronics, including LEDs and field-effect transistors. The work uses a combination of molecular dynamics and electronic structure techniques to study the effect of changes in the conformation of the molecules on their transport properties. A second article describes work on the bonding properties of novel compounds of uranium and nitrogen. These studies combine electronic structure calculations and experiment to make deductions about ligand exchange with these uranium compounds. A related article addresses the issue of carbon-hydrogen bond activation in actinide complexes. This work is important to the understanding of how these complexes can be used to synthesize a diverse array of actinide organometallic compounds with various applications.

In yet another article on organometallic complexes there is a description of electronic structure calculations designed for comparison with x-ray absorption spectroscopy experiments. This work investigates the nature of metal-chloride bonds and how they affect the properties of the compounds. Another article describes calculations designed to start understanding corrosion effects in uranium and transuranic metals. The chemistry that is studied here can have significant applications for

the storage of radioactive waste. The last article in the chemistry series studies organometallic compounds with potential applications in electronic devices. Here, a combination of theoretical and experimental techniques is employed to probe the spin degrees of freedom in porphyrins.

The articles on bioscience deal with the structural biology of amyloid oligomers associated with Alzheimer’s disease. These are typically portions of proteins that fold differently than in the normal protein. The work described here involves atomic-scale simulations to determine the way in which these Alzheimer’s-associated peptides fold. Another article in bioscience describes work on molecular epidemiology dealing with the identification of viral subtypes. The work employs statistical analysis to guide the classification of hepatitis C virus (HCV) and human immunodeficiency virus 1 (HIV-1) variants. A third article in the bioscience series deals with work on the dynamics of virus and immune response to experimental vaccines for HIV. At the end of the section we find an article on the structural conformation of beryllium compounds associated with chronic beryllium disease.



# Atomic Description of Alzheimer's Amyloid Beta Oligomers

S. Gnanakaran, T-10; Ruth Nussinov, National Cancer Institute;  
Angel E. Garcia, Rensselaer Polytechnic Institute

Cells have elaborate procedures to ensure that proteins fold correctly. In the event that a protein misfolds, cellular mechanisms are in place to detect and degrade it before it can become toxic. Despite these efforts, a range of debilitating human diseases is associated with protein misfolding. Lately, most attention has been paid to a group of diseases where proteins or peptides convert from their normally soluble forms to aggregates of insoluble fibrils or plaques. The final forms of these aggregates often have an ordered assembly of cross  $\beta$ -sheet fibrils and are referred to as amyloids. The deposits of amyloid are associated with at least 20 diseases, including such diverse entities as Alzheimer's disease (amyloid beta), prion diseases (prion), familial polyneuropathy (transthyretin), type 2 diabetes mellitus (islet amyloid or amylin), systemic amyloidosis (lysozyme or transthyretin), Parkinson's disease ( $\alpha$ -synuclein), and Huntington's disease (huntingtin). The fact that many of the ordinary proteins can form amyloid fibrils under conditions that bias the population of unfolded conformations makes the study of fibril formation very intriguing.

All-atom simulations have been carried out on monomer and dimer of the aggregation-prone fragment (16-22) of amyloid beta peptide ( $A\beta_{16-22}$ ), which is implicated in Alzheimer's disease (Fig. 1). It is motivated by the recent clinical studies that have suggested oligomers as the possible pathological agents. A thorough understanding of dimer formation provides a detailed picture of forces that drive the interaction between fragments. The theoretical challenge is to capture the relevant configurational ensembles of the  $A\beta_{16-22}$  dimers at the atomic level as a function of temperature (T) in explicit aqueous solution. The replica exchange molecular dynamics method, which has been successfully applied to peptide folding, is used as a means of sampling the configurational space with proper Boltzmann weighting so that the structural, motional, and thermodynamic description of self-assembly can be obtained.

The free-energy landscape showing the delicate balance between different monomer and dimer conformations is mapped along carefully chosen reaction coordinates (Fig. 2). The canonical ensembles at 38 different T are used to describe the thermodynamics and the relative stabilities of at least six different

dimer conformations, including parallel and antiparallel orientations. In summary, this all-atom simulation study reveals that dimers of the aggregation-prone fragment of  $A\beta$  peptide do not necessarily adopt only the parallel and antiparallel conformations commonly seen in the amyloid fibril. Six highly populated conformations are found, with the parallel conformation preferred at lower T, and the antiparallel conformation is preferred at higher T (Fig. 3). Furthermore, we delineate the nature of the molecular forces that activate and stabilize these different dimer conformations as a function of T, especially as related to the secondary structural propensity of the monomer [1]. This study also identified parallel loop dimer conformations that are stabilized due to specific interactions with water molecules. Currently, aggregation of a large number of fragments of the same peptide is under consideration (Fig. 4). Calculations were carried out on Pink as part of the Institutional Computing Program.

For more information contact S. Gnanakaran at [gnana@lanl.gov](mailto:gnana@lanl.gov).

[1] S. Gnanakaran et al., *J Am Chem Soc*, **128**, 2158 (2006).

## Funding Acknowledgments

- Los Alamos National Laboratory Directed Research and Development Program




Fig. 1. Aggregation prone segment of amyloid beta peptide.




Fig. 2. Free energy landscape of  $A\beta_{16-22}$  dimers at 310K. Representative conformations of dimers are shown.




Fig. 3. Temperature profile of dimers identified in Fig. 2.




Fig. 4. Snapshot from simulations of six peptide fragments.

# Automatic Classification of HCV and HIV-1 Sequences with the Branching Index

Peter Hraber, Carla Kuiken, Mark Waugh, Shaun Geer, William J. Bruno, Thomas Leitner, T-10

**K**nowledge of viral subtype is important for molecular epidemiology, for identification of new variants and circulating recombinant forms, and for informed treatment options. Viral sequences are tremendously variable, presenting a challenge for systematic classification.

Automatic classification of viral sequences should be fast, objective, accurate, and reproducible. Most methods that classify sequences use either pairwise distances or phylogenetic relations. However, they cannot discern when a sequence is unclassifiable. The branching index (BI) combines distance and phylogeny methods to compute a ratio that quantifies how closely a query sequence clusters with a subtype clade (Fig. 1).

In the hypothesis-testing framework of statistical inference, the BI is compared with a threshold to test whether sufficient evidence exists for the query sequence to be classified among known sequences. If above the threshold, the null hypothesis of no support for the subtype relation is rejected, and the sequence is taken as belonging to the subtype clade with which it clusters on the tree.

We studied statistical properties of the BI for subtype classification in HCV and HIV-1. Pairs of BI values with known positive and negative test results (Fig. 2) were computed from 10,000 random fragments of reference alignments. Sampled fragments were of sufficient length to contain phylogenetic signal that properly groups reference sequences together into subtype clades. For HCV, a threshold BI of 0.71 yields 95.1% agreement with reference subtypes, with equal false positive and false negative rates. For HIV-1, a threshold of 0.66 yields 93.5% agreement. Higher thresholds can be used where lower false positive rates are required.


BI profiles result from moving windows over the length of the query sequence. Shown graphically, BI profiles enable visualization of classification consistency over the query sequence (Fig. 3). A new service for automatic classification of HIV-1 and HCV sequences with the branching index is being provided online.

This classification approach might also be used to classify influenza and other viral genotypes, to identify bacterial subspecies, for molecular subtyping of bacterial toxins, and for more general use in clustering algorithms and supervised learning.

For further information contact Peter Hraber at [phraber@lanl.gov](mailto:phraber@lanl.gov).

## Funding Acknowledgements

- National Institutes of Health
- Department of Energy



*Fig. 1. BI illustration. The branching index is defined for query sequence  $x$  and subtype clades  $S_{1-4}$  as the ratio of branch lengths  $\lambda_a / (\lambda_a + \lambda_b)$ , given proximity to  $S_3$ .*

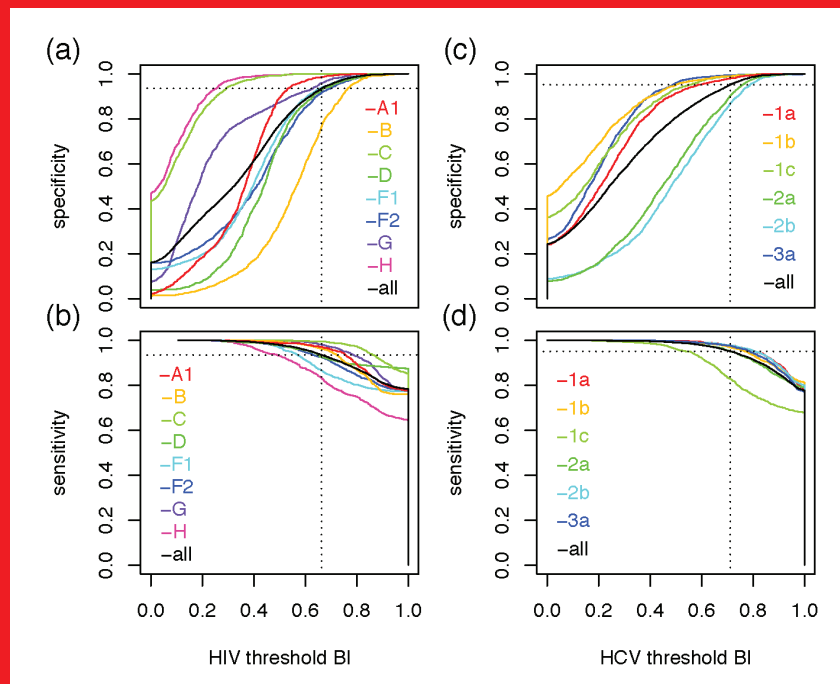


Fig. 2. Error rates from tests with known outcomes. Cumulative distributions from 10,000 random HIV-1 fragments show specificity (a) and sensitivity (b) for any threshold, and likewise for HCV specificity (c) and sensitivity (d). Line color indicates viral subtype. Dotted lines indicate thresholds that optimize classification accuracy for all samples.




Fig. 3. BI profile from HCV sequence AY651061 (400 nt windows with 40 nt between each successive window) indicates a 1a/1c recombinant.

# Recent Vaccination Improves the Early Control of Virus in SHIV Infection but not the Clinical Outcome

Leslie L. Chavez, Alan S. Perelson, Ruy M. Ribeiro, T-10; Miles P. Davenport, Univ. of New South Wales; John W. Shiver, Lynda G. Tussey, Kara S. Cox, Margaret Bachinsky, Fubao Wang, Lingyi Huang, William A. Schleif, Mary-Ellen Davies, Aimin Tang, Danilo R. Casimiro, Merck Research Laboratories

The United Nations Acquired Immunodeficiency Syndrome (UNAIDS) organization reported that, at the end of 2007, approximately 40 million people were living with human immunodeficiency virus (HIV) [1]. Of these, 25.8 million (64%) were located in the developing countries of sub-Saharan Africa (Fig. 1). Since the beginning of the epidemic over 25 million people have died of the complications occurring in late stage HIV [1], when acquired immunodeficiency syndrome (AIDS) manifests itself. The best hope the world has to control this epidemic is the development of efficient vaccines.

We have been working to understand the dynamics of virus and immune response within the host for different experimental vaccines. How do they affect the early viral growth kinetics [2]? What are the dynamics of the immune response [3-5]? We have been collaborating with Merck & Co, Inc., to analyze its experimental data on a CD8<sup>+</sup> T-cell inducing vaccine. Our objective was to shed light on the effects of the time that elapses between vaccination and infection on the performance of the vaccine.

Twenty-eight rhesus macaques were given a vaccine-boost protocol, and groups of four macaques were then challenged with virus at 1, 3, 6, 12, and 24 weeks after the boost. We used statistical methods to characterize the kinetics of viral load, CD4<sup>+</sup> and virus-specific CD8<sup>+</sup> T-cells. We found that all vaccinated animals showed an increase in CD8<sup>+</sup> T-cells compared to control animals, but that the number of virus-specific CD8<sup>+</sup> T-cells had an exponential decay between 1 and 12 weeks following vaccination (Fig. 2). However, the viral and T-cell kinetics over the first 2 weeks differed between the vaccinated groups (Fig. 3), with more recent vaccination improving the early control of virus. Interestingly, the rates of virus-specific CD8<sup>+</sup> T-cell expansion were greater in animals having higher viral loads at 1 week.

Differences in time of challenge since vaccination did not lead to differences in the clinical outcome for vaccinated animals. However, the dynamics of the virus after recent vaccination was qualitatively different from later challenges and these dynamics of early viral load seems to have a role in virus-specific CD8<sup>+</sup> T-cell generation.

These studies are important in thinking about what will be the correct protocol for vaccination, boost, and revaccination once an effective HIV vaccine is developed.

For further information contact Ruy M. Ribeiro at [ruy@lanl.gov](mailto:ruy@lanl.gov).

- [1] UNAIDS/WHO AIDS Epidemic Update, UNAIDS/WHO, Geneva (2007).
- [2] M.P. Davenport, et al., *J Virol*, **78**, 10096-10103 (2004).
- [3] M.P. Davenport, et al., *J Virol*, **79**, 10059-10062 (2005).
- [4] M.P. Davenport, et al., *J Acquir Immune Defic Syndr*, **41**, 259-265 (2006).
- [5] M.P. Davenport, et al., *Immunol Rev*, **216**, 164-175 (2007).

## Funding Acknowledgments

- James S. McDonnell Foundation
- National Institutes of Health


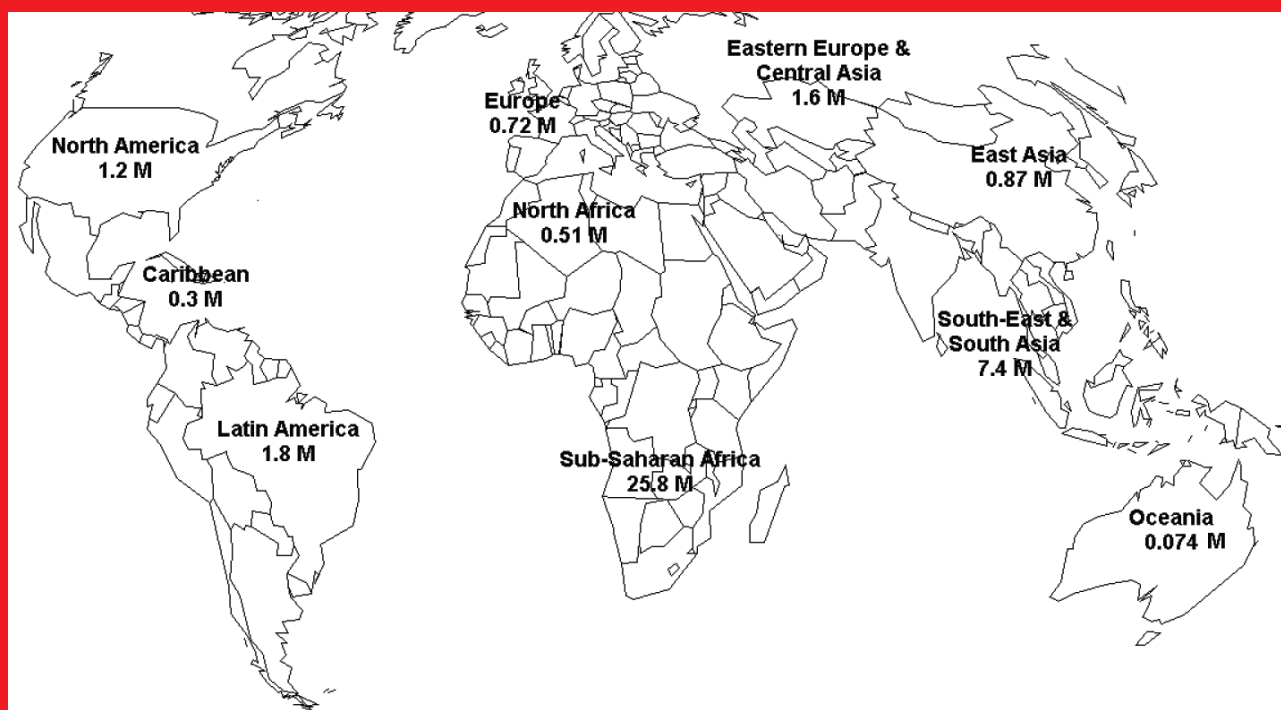



Fig. 2. Decay and plateau of virus-specific CD8<sup>+</sup> T-cells after vaccination and before viral challenge. Individual animals (colored circles) and group median (black diamonds) are shown. In the vaccinated and boosted groups, the animals' virus-specific CD8<sup>+</sup> T-cells drop monotonically with time between weeks 1 and 12.





*Fig. 1. The burden of HIV. Approximate number and geographical distribution of HIV-infected people in the world.*



*Fig. 3. Acute infection profile of virus-specific CD8<sup>+</sup> T-cells differs by challenge time (profiles are averaged over the four macaques in each group).*

# Towards a Theory of Transport in Disordered Organic Polymers: Poly(*p*-Phenylene Vinylene) (PPV)

Ping Yang, Enrique R. Batista, Sergei Tretiak, Richard L. Martin, T-12; Avadh Saxena, Darryl L. Smith, T-11

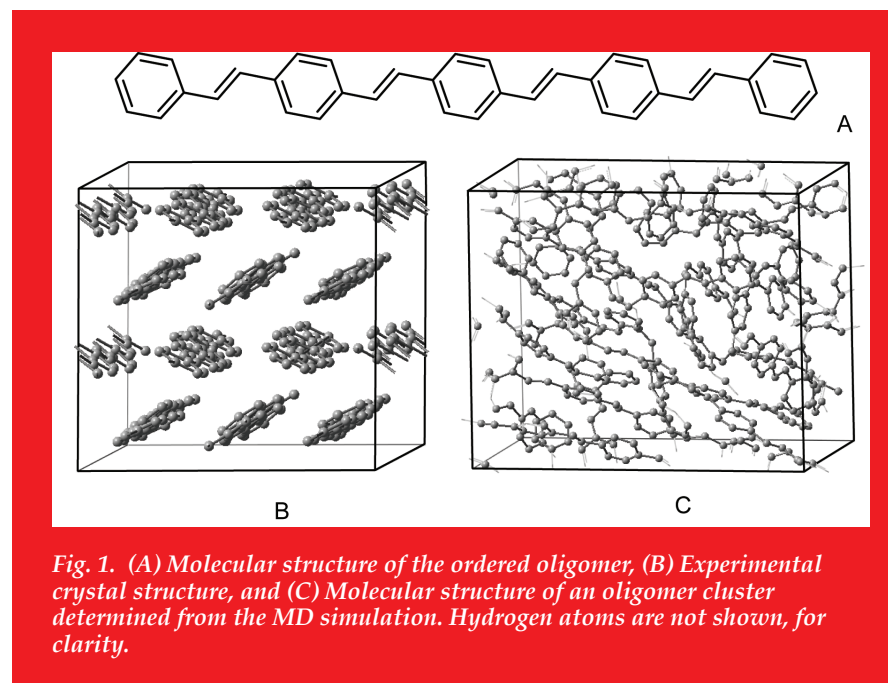
The emergence of organic electronic devices, including light-emitting diodes and field-effect transistors, fabricated from conjugated polymers such as poly(*p*-phenylene vinylene) (PPV), has stimulated research into the electrical and optical properties of these semiconducting polymers. The nature of the charge-trapping states that stem from the conformational disorder in these systems has not been clearly identified from a structural perspective. A complete description of these electronic materials is challenging because they are highly disordered and interactions between molecules are very important in determining their properties. Despite extensive experimental and theoretical research, there is still a limited understanding of these strongly interacting disordered materials.

Our approach uses a combination of classical molecular dynamics (MD) and density functional theory (DFT) to investigate the role of intra-molecular conformational disorder and inter-molecular electronic interactions on the electronic structure of disordered clusters of PPV. Classical molecular dynamics is used to determine probable molecular geometries for the interacting disordered oligomers (Fig. 1). For an ensemble of such configurations, the electronic structure and associated energy levels are determined with first-principles density functional theory (Fig. 2). Intra-molecular and inter-molecular effects are disentangled by contrasting results for densely packed oligomer clusters with those for ensembles of isolated oligomers with the same intra-molecular geometries (Fig. 3). We find that in PPV, electron trap states are induced primarily by intra-molecular configuration disorder, while the hole trap states are generated primarily from inter-molecular electronic interactions. The molecular orbitals responsible for the deep trap states can be examined (Fig. 4), and associated with specific conformational properties.

For further information contact Ping Yang at [pyang@lanl.gov](mailto:pyang@lanl.gov).

## Funding Acknowledgments

- Department of Energy, Office of Science, Office of Basic Energy Sciences



*Fig. 1. (A) Molecular structure of the ordered oligomer, (B) Experimental crystal structure, and (C) Molecular structure of an oligomer cluster determined from the MD simulation. Hydrogen atoms are not shown, for clarity.*




Fig. 2. Energy level diagrams from the DFT calculations. The upper panel shows results for the ideal crystal (left) and 14 disordered oligomer clusters whose geometry was determined from the MD calculations; the lower panel shows results for an isolated ordered oligomer (left) and an ensemble of 12 isolated oligomers with the same molecular geometries as in the corresponding column of the upper panel.




Fig. 3. Calculated density of states for: (A) the ideal crystal, (B) disordered cluster S4, (C) disordered cluster S5, (D) the ensemble average of clusters S1-S14, (E) the isolated oligomer in the ideal crystal, (F) the 12 isolated oligomers in cluster S4, (G) the 12 isolated oligomers in cluster S5, and (H) the ensemble average of isolated oligomers in clusters S1-S14. Note that green lines denote different energy gaps in top panels A-D versus bottom panels E-H.

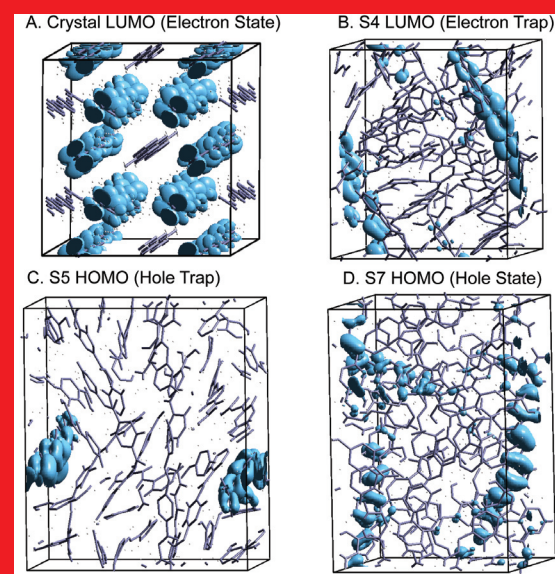


Fig. 4. Calculated electron densities for selected states: (A) the crystal lowest unoccupied molecular orbital (LUMO) state, (B) the LUMO level for cluster S4, (C) the highest occupied molecular orbital (HOMO) level for cluster S5, and (D) the HOMO level for cluster S7.

# Theoretical Study of Ligand Exchange Reaction on Bis(Imido) Uranium(VI)

Ping Yang, Enrique R. Batista, T-12; Liam P. Spencer, Brian L. Scott, James M. Boncella, MPA-MC

**F**or the past 150 years, studies of uranium(VI) have been generally directed towards understanding the chemical behavior and unique bonding in the uranyl ion ( $\text{UO}_2^{2+}$ ). This ion possesses U-O bonds that have high thermodynamic stability and extreme kinetic inertness. As a result, the majority of  $\text{UO}_2^{2+}$  reaction chemistry involves substitution of equatorially coordinated ligands while leaving the U-O bond unaffected [1]. Recently, we have reported the syntheses of isoelectronic bis(imido)  $[\text{U}(\text{NR})_2]^{2+}$  and oxo-imido  $[\text{U}(\text{NR})(\text{O})]^{2+}$  ions that possess many of the bonding features found in  $\text{UO}_2^{2+}$  [2].

Here we report investigations of the reactivity of the  $[\text{U}(\text{N}^t\text{Bu})_2]^{2+}$  (Fig. 1.) ion towards aryl isocyanates. Given the strength and stability of the U=O multiple bond, these reactions were anticipated to produce the oxo-imido framework. Instead, we observed an unexpected reaction that does not involve U=O bond formation, but rather an imido exchange reaction in which aryl-imido for alkyl-imido substitution occurs.

Stirring orange-red solutions of **1** with PhNCO in  $\text{CH}_2\text{Cl}_2$  generates dark red-brown solutions. In the reaction with PhNCO, the mixed bis(imido) complex,  $\text{U}(\text{NPh})(\text{N}^t\text{Bu})(\text{I})_2(\text{OPPh}_3)_2$  (**3**), can be isolated from this reaction, a complex that can be well characterized via  $^1\text{H}$  NMR and crystallography (Fig. 2). While there is precedence for this aryl- for alkyl-imido substitution, it is not clear why this reaction does not form a thermodynamically and kinetically stable U=O bond. To provide some insight on why this exchange does not produce a U=O bond, density functional theory (DFT) calculations were performed to elucidate the relative energies of the intermediates and products of this reaction. There are two reaction pathways which could generate the bis(imido) product **3** and the oxo-imido species **2** from  $\text{U}(\text{N}^t\text{Bu})_2(\text{I})_2(\text{OPPh}_3)_2$  (Paths 1 and 2, Fig. 1). The first pathway involves the [2+2] cycloaddition of the C=N bond of the aryl isocyanate to form an N,N-bound ureato intermediate (**7**), which can isomerize to form species **8** with the -NPh group trans to the *tert*-butyl imido moiety. Compound **8** can then eliminate

*tert*-butyl isocyanate to generate the unsymmetrical bis(imido) complex **3** (Path 1). Alternatively, N,O-bound carbamate intermediates **9** and **10** could form, which result from the [2+2] cycloaddition of the C=O bond of the aryl isocyanate across the U=N imido bond. Elimination of a substituted carbodiimide would generate the oxo-imido complex **2** (Path 2).

The computational results suggest the lowest energy pathway involves the [2+2] cycloaddition of the C=N bond of phenyl isocyanate to form the N,N-bound ureato intermediate **7** (Path 1, Fig. 3). In the calculations performed, it is assumed that  $\text{OPPh}_3$  dissociation occurs in order to generate intermediates **7** and **8**. Experimentally, the reactions between **1** and PhNCO proceed much more slowly in the presence of excess  $\text{OPPh}_3$ , as is consistent with this assumption. Overall the transformation of  $\text{U}(\text{N}^t\text{Bu})_2(\text{I})_2(\text{OPPh}_3)_2$  (**1**) to  $\text{U}(\text{NPh})(\text{N}^t\text{Bu})(\text{I})_2(\text{OPPh}_3)_2$  (**3**) is exergonic, with the free energy of the bis(*tert*-butyl)imido uranium complex (**1**) + PhNCO higher than the mixed imido species **3** +  $^t\text{BuNCO}$  by 5.6 kcal/mol.

In contrast to this mechanism, the calculated [2+2] C=O cycloaddition bond pathway (Path 2) involves the formation of the higher energy N,O-bound ureato intermediates **9** ( $19.8 \text{ kcal mol}^{-1}$ ) and **10** ( $28.7 \text{ kcal mol}^{-1}$ ). Complex **9** can then eliminate the mixed carbodiimide  $^t\text{BuN}=\text{C}=\text{NPh}$  and form the oxo-imido species **2**. As anticipated, the relative energy of **2** is substantially lower in energy than the unsymmetrical imido species **3** ( $15.2 \text{ kcal/mol}$ ). This energy difference between U=O and U=N bond formation has also been observed in cyclopentadienyl-substituted uranium(IV) complexes.

Given the difference in the relative energies of the intermediates between these two pathways and the results from  $^{15}\text{N}$ -labeling studies, it appears the mechanism for the formation of **3** involves the [2+2] cycloaddition of the aryl isocyanate C=N bond across the U=N imido moiety. These results are quite surprising given the thermodynamic and kinetic stability of U=O bonds. However, the experiment repeated with labeled  $^{15}\text{N}$  supports the proposed exchange reaction mechanism.

**For further information contact Enrique R. Batista at [erb@lanl.gov](mailto:erb@lanl.gov).**

[1] J.J.M. Katz and G.T. Seaborg, in *The Chemistry of the Actinide Elements*, 2<sup>nd</sup> ed., Chapman and Hall: London, 1137-1140 (1986).

[2] T.W. Hayton et al., *Science*, **310**, 1941 (2005).

## Funding Acknowledgements

- Department of Energy, Office of Science, Office of Basic Energy Sciences
- Glenn T. Seaborg Institute




Fig. 1. Molecular structure of  $[U(NPh)_2(N^tBu)I_2(OPPh_3)_2](1)$ :  $U-N=1.850\text{\AA}$ ,  $U-O=2.408\text{\AA}$ ,  $UI=3.169\text{\AA}$ ,  $O-P=1.528\text{\AA}$ ,  $N-U-N=175\text{ deg}$ .




Fig. 2. Given the strength of  $U=O$  bonds, it was anticipated that the reaction of the bis(imido) complex  $U(N^tBu)_2I_2(OPPh_3)_2$  (11) with aryl isocyanates would yield the oxo-imido complex 2 or possibly an  $N,O$ -bound ureate.

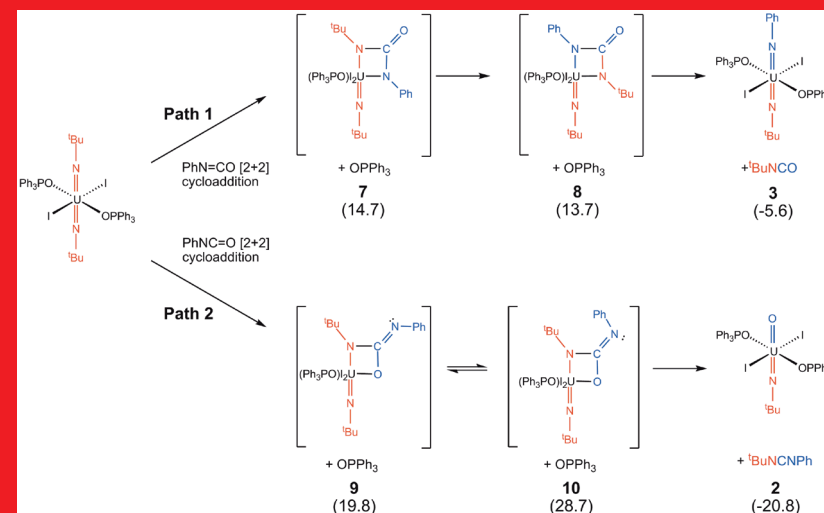


Fig. 3. Relative free energies of the products and intermediates in the potential pathways for the formation of 2 and 3. Energies reported at the hybrid DFT level of theory relative to the energy of 1 are provided in parentheses in kcal/mol.



# Theoretical Studies of Carbon-hydrogen Bond Activation Chemistry of Actinide Complexes

Ping Yang, Ingolf Warnke, Richard L. Martin, P. Jeffrey Hay, T-12


The thorium(IV) and uranium(IV) alkyl complexes  $(C_5Me_5)_2AnR_2$  (where An = Th, U; R =  $CH_3$ , CPh, Ph) have proven to be versatile starting materials for the synthesis of a diverse array of actinide organometallic systems containing An-N bonds such as imido, hydrazonato, and ketimido complexes, which feature novel electronic properties [1]. Recently, the Kiplinger group at Los Alamos National Laboratory in MPA-10 reported that these actinide alkyl complexes undergo interesting C-H and C-N bond cleavage chemistry with N-heterocycles [2] such as 2-picoline (2-methylpyridine), which possesses both  $sp^2$  and  $sp^3$  hybridized C-H bonds. The thorium alkyl complex  $(C_5Me_5)_2Th(CH_3)_2$  activates both the  $sp^3$  C-H bond on the 2-picoline methyl group and the *ortho*  $sp^2$  C-H bond on the ring. This is in marked contrast to the uranium system, which only reacts with a  $sp^2$  C-H bond on the 2-picoline aromatic ring (Fig. 1). Deuterium labeling studies demonstrated that the thorium and uranium  $(C_5Me_5)_2An(CH_3)_2$  complexes react with 2-picoline by different mechanistic reaction pathways.

Density functional theory (DFT) methods have been employed to explore actinide-ligand interactions, such as structures, thermochemistry, and spectroscopic properties, using hybrid functionals [3,4]. We have performed a computational study of the competitive  $sp^2$  versus  $sp^3$  C-H activations. The products and resulting thermochemistry in these reactions are compared, and likely reaction precursors and transition states are also identified, as well the plausible reaction pathways.

The results of theoretical study are consistent with reported experimental results. Optimized geometries are in excellent agreement with X-ray crystal data. The calculated reaction energies prove that the  $sp^2$  C-H bond activation product is the most stable structure (i.e., thermodynamic) product. Both theoretical and experimental observations point to the  $sp^3$  product as the kinetic product and the  $sp^2$  product as the thermodynamic product for thorium, while the  $sp^2$  product is both the kinetic and thermodynamic selection in the uranium scheme

(Fig. 1). The reaction initiates from formation of a weakly bound adduct, followed by the activation of adjacent C-H activation by An-center leading to an agostic transition state. We find that the actinide atom plays a fundamental role during the hydrogen migration process from 2-picoline to the methyl leaving group. Agostic 5-centered transition structures for the actinide C-H activation reaction pathways have not been reported before, to the best of our knowledge. The origin of the regioselectivity of these reactions rests in these highly ordered configurations of transition states. Despite many common features found between thorium and uranium systems, including the similar geometries of products, adducts and the agostic transition states, the relative activation energies between  $sp^2$  and  $sp^3$  activation differ slightly:

$$\text{Th: } \Delta E_{sp^3}^\ddagger < \Delta E_{sp^2}^\ddagger \quad \text{U: } \Delta E_{sp^2}^\ddagger < \Delta E_{sp^3}^\ddagger$$



On the basis of the combination of labeling, structural, and computational information, we proposed a general mechanism for the C-H activation of N-heterocycle by actinocene complexes. “Agostic migration” cyclometalation is indicated as an operative mechanism.

Table 1: The calculated 3D structures of the transition states involved in the  $sp^2$  and  $sp^3$  C-H bond activation of 2-picoline from the most stable thorium and uranium adducts.

	$sp^2$ activation	$sp^3$ activation
Complex		
Thorium		
Uranium		

For further information contact P. Jeffrey Hay at [pjhay@lanl.gov](mailto:pjhay@lanl.gov).

- [1] J. A. Pool, B.L. Scott, and J.L. Kiplinger, *J. Am. Chem. Soc.*, **127**, 1338-1339 (2005).  
 [2] J. L. Kiplinger et al., *J. Alloys Compd.*, **444-445**, 477-482 (2007).  
 [3] N. Kaltsoyannis, *Chem. Soc. Rev.*, **32**, 9-16 (2003).  
 [4] N. Kaltsoyannis et al., Theoretical Studies of the Electronic Structure of Compounds of the Actinide Elements in *The Chemistry of the Actinide and Transactinide Elements*, Vol. 3, 3 ed., Springer, Dordrecht, The Netherlands (2006).

#### Funding Acknowledgments

- Department of Energy, Office of Science, Office of Basic Energy Sciences
- Glenn T. Seaborg Institute

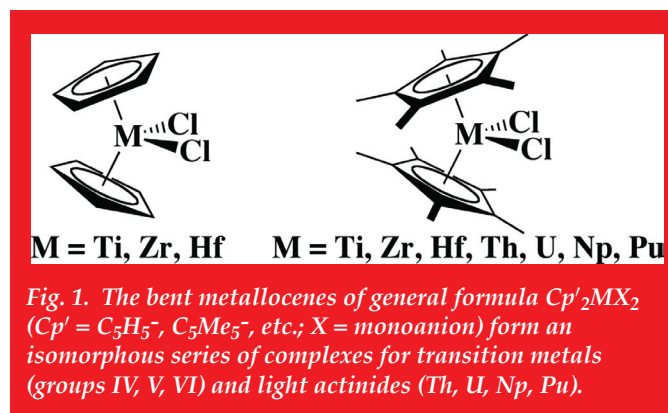
# Covalency Trends in Group IV Metallocene Dichlorides: Chlorine K-edge X-ray Absorption Spectroscopy and Time Dependent-density Functional Theory

Ping Yang, Enrique R. Batista, P. Jeffrey Hay, Richard L. Martin, T-12; Stosh A. Kozimor, Carol J. Burns, C-NR; Kevin S. Boland, Laura E. Wolfsberg, C-IIAC; David L. Clark, Daniel E. Schwarz, ADSMS; Steven D. Conradson, Juan S. Lezama, MST-8; Marianne P. Wilkerson, C-ADI; Christin N. Christensen, New Mexico State Univ.

The bent metallocenes of general formula  $\text{Cp}'_2\text{MX}_2$  ( $\text{Cp}' = \text{C}_5\text{H}_5^-, \text{C}_5\text{Me}_5^-, \text{etc.}$ ;  $\text{X} = \text{monoanion}$ ) form an isomorphous series of complexes for transition metals (groups IV, V, VI) and light actinides (Th, U, Np, Pu), and are among the most extensively utilized molecular systems in organometallic research. Theoretical and spectroscopic studies of metallocenes indicate that variations in orbital character and metal  $d$ - and  $f$ -covalency are connected with differences in chemical reactivity. Of the experimental approaches to determine covalency, ligand K-edge x-ray absorption spectroscopy (XAS) has emerged as a powerful measure of the amount of ligand orbital mixing with metal-based orbitals. On the low-energy side of the ligand K-edge are bound state transitions of ligand  $1s$  electrons into metal-based orbitals that contain ligand  $p$  character due to covalent mixing. The intensity of these pre-edge features provides a direct experimental measure of covalency.

Since the  $\text{Cp}'_2\text{MCl}_2$  compounds form an isomorphous series for transition metals and light actinides (Fig. 1), they are well-suited for a ligand K-edge XAS study to compare relative changes in orbital mixing between a common  $\text{Cl}^-$  ligand and  $3d$ ,  $4d$ ,  $5d$ , and  $6d/5f$  orbitals in metal-ligand bonding. Moreover, their electronic structure is well understood, and the Cl K-edge XAS of the first member of this series,  $\text{Cp}_2\text{TiCl}_2$  ( $\text{Cp} = \text{C}_5\text{H}_5^-$ ), has been recently analyzed by both experiment and theory, thereby providing a standard for comparison with heavier congeners.

For a more detailed understanding of M-Cl bonding interactions we employed electronic structure calculations on  $\text{Cp}_2\text{MCl}_2$  complexes using B3LYP hybrid density functional theory (DFT) in the *Gaussian 03* code. The Stuttgart 97 relativistic effective core potential and associated basis sets (minus the most diffuse function) were used for Ti, Zr, and Hf, and the 6-31 G\* basis sets for C and H.



The first level of theory assumes that the transition amplitudes are directly proportional to the % Cl  $p$  character of the virtual orbitals. A more rigorous level of theory involves a linear response theory (TD-DFT)

calculation, where the probability amplitudes were extracted from the transition dipole moments between the calculated excited states and the ground states. The excitations originating from all the intermediate states between the Cl  $1s$  and the highest occupied molecular orbital (HOMO) were not included so that only excitations from the core levels to virtual molecular orbitals (MOs) could be analyzed. This allows the virtual orbitals to mix among themselves to reflect the presence of the core hole in Cl. However, we do not include relaxations in the occupied orbitals associated with the core hole.

The electronic structure of bent metallocene dichlorides is well understood. The bent  $\text{Cp}_2\text{M}$  fragment is characterized by six M-Cp orbital interactions at lower energy, leaving three low-energy metal orbitals ( $1a_1$ ,  $1b_2$ , and  $2a_1$ ) and two higher-lying metal orbitals ( $1b_1$ ,  $2b_2$ ) available to form bonding interactions with the Cl ligands. While the first set of three orbitals is generally regarded as the most important for M-Cl bonding, all five are necessary to understand the XAS intensities. The five resulting M-Cl antibonding orbitals are shown in Fig. 2.

From ground state calculations, the total calculated % Cl  $3p$  character for the metal-based orbitals is 20%, 18%, and 17%,  $\text{Cp}_2\text{MCl}_2$  complexes ( $\text{M} = \text{Ti, Zr, and Hf}$ ; **1**, **2**, and **3**, respectively), which agrees well with the experimental values. From the more rigorous TD-DFT calculations we were able to simulate the Cl K-edge XAS, and this is compared with experimental XAS in Fig. 3. The TD-DFT calculations reproduce the basic experimental features of peak intensity, peak area, and peak energies remarkably well, and provide additional confidence on spectral assignments. The calculated and experimental spectra for **1-3**, Fig. 3, show a decrease in amplitudes and increase in peak width for the second pre-edge feature. The increase in width stems from the larger splitting among the  $d$  virtual states, and the decrease in intensity to an overall reduction in oscillator strength.




Fig. 2. Shape of five M-Cl antibonding orbitals in the electronic structures of bent metallocene dichlorides  $\text{Cp}_2\text{MX}_2$  ( $\text{M} = \text{Ti}, \text{Zr}, \text{Hf}$ ).

For all three compounds, our TD-DFT calculations suggest that the two lowest unoccupied orbitals are the  $1a_1$  and  $1b_2$  M-Cl anti-bonding orbitals. Hence, for **1-3** the first XAS pre-edge peak is attributed to a  $\text{Cl } 1s \rightarrow 1a_1$  transition. While it is tempting to assign the second pre-edge peak to a  $1s \rightarrow 1b_2$  transition, the calculations reveal that the  $1b_2$  orbital is pushed sufficiently high in energy such that the manifold of four primarily metal nd molecular orbitals ( $1b_2$ ,  $1b_1$ ,  $2b_2$ , and  $2a_1$ ) are closely grouped together, spanning only 0.6, 0.9, and 1.1 eV for **1-3**,




Fig. 3. The Cl K-edge experimental data (blue), the TD-DFT calculated spectrum (red dashes), and the TD-DFT calculated transitions (black bars) for polystyrene films of  $\text{Cp}_2\text{MCl}_2$  ( $\text{M} = \text{Ti}$  top,  $\text{Zr}$  middle, and  $\text{Hf}$  bottom; **1**, **2**, and **3**, respectively).

respectively. Depending on the metal ion, all four metal d orbital mix with the Cl  $3p$  orbitals to varied extents with orbital mixing following the order  $\text{Hf} > \text{Zr} > \text{Ti}$ . For example, in **1** the  $1s \rightarrow 1b_2$  transition contributes 87% to the total intensity of the second peak, while smaller contributions from the  $1s \rightarrow 1b_1$ ,  $2b_2$ , and  $2a_1$  transitions provide the additional 13%. In contrast, the amount of metal Cl  $3p/\text{Hf } 5d$  mixing is more evenly distributed in **3** and the  $1s \rightarrow 1b_2$  transition accounts for only 56% of the total intensity of the second peak. These calculations are consistent with the experimental spectra for **1-3**, which show decreased amplitudes and increased peak widths for the second pre-edge features.

Both theory and experiment indicate that the M-Cl bond for the  $3d$  complex, **1**, is more covalent than the  $5d$  analogue, **3**. Although the covalency in the Zr-Cl bond for **2** was experimentally indistinguishable from that of the Ti-Cl bond in **1**, the calculations suggest higher % Cl  $3p$  character for **1** than for **2**. Since the percent of Cl  $3p$  character in the  $1a_1$  molecular orbital was calculated, as well as found experimentally, to be the same for all three complexes, the differences in total covalency for the M-Cl bonds are attributable to the degree of Cl  $3p$  orbital mixing with the  $1b_2$ ,  $1b_1$ ,  $2b_2$ , and  $2a_1$  manifold of metal d orbitals, which decreases from **1** to **3**. This finding is similar to the trends revealed by gas phase photoelectron spectroscopy of occupied orbitals in metallocene complexes reported by Marks and coworkers.

Theoretical and spectroscopic studies of metallocenes indicate that variations in orbital character and metal d-covalency are connected with differences in chemical reactivity. Our results demonstrate the utility of Cl K-edge XAS for quantitatively analyzing the covalent mixing in M-Cl bonds along the series of  $3d$ ,  $4d$ , and  $5d$  group IV complexes. Along this series, covalency decreases, and the % Cl character is concomitantly spread over a greater number of MOs as  $Z$  increases. These data also demonstrate that reliable intensity data can be obtained using polystyrene-encapsulated samples, which should allow us to extend these studies to the heavier actinide series where both  $5f$  and  $6d$  orbitals are expected to play a role in M-Cl bonding, and where transition assignments must rely on accurate theoretical calculations.

These studies have now been extended to the light actides, where evidence of covalency in the  $5f$  orbitals has been observed. The degree of interaction is roughly half that observed in the transition metal benchmarks. Further experimental and theoretical studies should reveal a great deal about this fundamental issue in chemical bonding.

For further information contact Richard L. Martin at [rlmartin@lanl.gov](mailto:rlmartin@lanl.gov).

#### Funding Acknowledgments

- Department of Energy, Office of Science, Office of Basic Energy Sciences
- Frederick Reines Postdoctoral Fellowship
- Glenn T. Seaborg Institute

## Spin-orbit Coupling Effects on the Electronic Structure of $\text{UO}_2$

Lindsay E. Roy, Richard L. Martin, T-12; Tomasz Durakiewicz, John J. Joyce, MPA-10; Gustavo Scuseria, Rice Univ.

**T**he corrosion reactions of uranium and transuranic metals pose a serious concern for environmental preservation. The long-term storage of radioactive waste (containing mainly uranium and plutonium compounds) requires a comprehensive knowledge of the possible oxidation reactions that could occur. These materials also represent a challenge for theorists; materials belonging to this class usually have partially occupied  $f$  orbitals where the electrons compete between itinerancy and localization.

When localized, the  $f$  electrons are strongly correlated, and these interactions normally lead to Mott-Hubbard insulators. For example, plutonium crystallizes in seven distinct phases, is highly reactive, and can have five different oxidation states when forming compounds. The surface of metallic plutonium easily oxidizes to  $\text{PuO}_2$  when exposed to air and moisture, while  $\text{Pu}_2\text{O}_3$  plays an important role in oxidation kinetics.

Given the incomplete understanding of the Pu oxidation chemistry and the myriad problems associated with experimental studies on these materials, theory can provide fundamental insight into the electronic structure and properties of these systems. Unfortunately, theoretical studies on these materials are quite difficult. The  $5f$  electrons can either be localized or contribute to bonding, and their relativistic effects and electron-electron correlations are very important factors in deciding the degree of localization. In particular, actinide oxides ( $\text{AnO}_2$ ) exhibit amazingly complex behavior, despite having a simple binary formula, because they straddle the metallic, ionic, and covalent bonding descriptions used by chemists and physicists. Most of these oxides are small band gap insulators and some exhibit antiferromagnetic ordering, or have uncertain magnetic ordering, at low temperature. Because of the strong correlations, previous theoretical studies on these materials using local spin-density approximation (LSDA) and generalized gradient approximations (GGA) of density functional theory (DFT) fail to accurately describe the bonding and electronic structure of these compounds.

Recently, we have shown that a third generation of functionals, the hybrid DFT approximation that combines the exact, non-local, Hartree-Fock exchange interaction with the traditional local (LDA) or semi-local GGA exchange and correlation interactions, is able to correctly predict a magnetic ground state, the insulating gap, and the lattice constants for  $\text{UO}_2$  series. In addition, the  $5f$  orbitals are nearly localized in the band structure of  $\text{UO}_2$ , but there is a small dispersion of  $\sim 300$  meV. This dispersion has recently been verified by colleagues in MPA-10.

We have recently included the effect of spin-orbit coupling (SOC), or the interaction of the electron spin magnetic moment with the magnetic moment due to the orbital motion of the electron, in our calculations. Our work on  $\text{UO}_2$  with SOC shows that the band gap decreases by only 0.06 eV, and that SOC is a minor perturbation. We are currently working on including SOC effects in the entire  $\text{AnO}_2$  series.

**For further information contact Lindsay E. Roy at [lroy@lanl.gov](mailto:lroy@lanl.gov).**

### Funding Acknowledgments

- Department of Energy, Office of Science, Office of Basic Energy Sciences






Fig. 1. The  $\text{CaF}_2$ -like lattice of actinide dioxides. The actinide and oxygen atoms form face-centered and simple cubic sublattices, respectively. Actinide and oxygen sites are in blue and red, respectively.

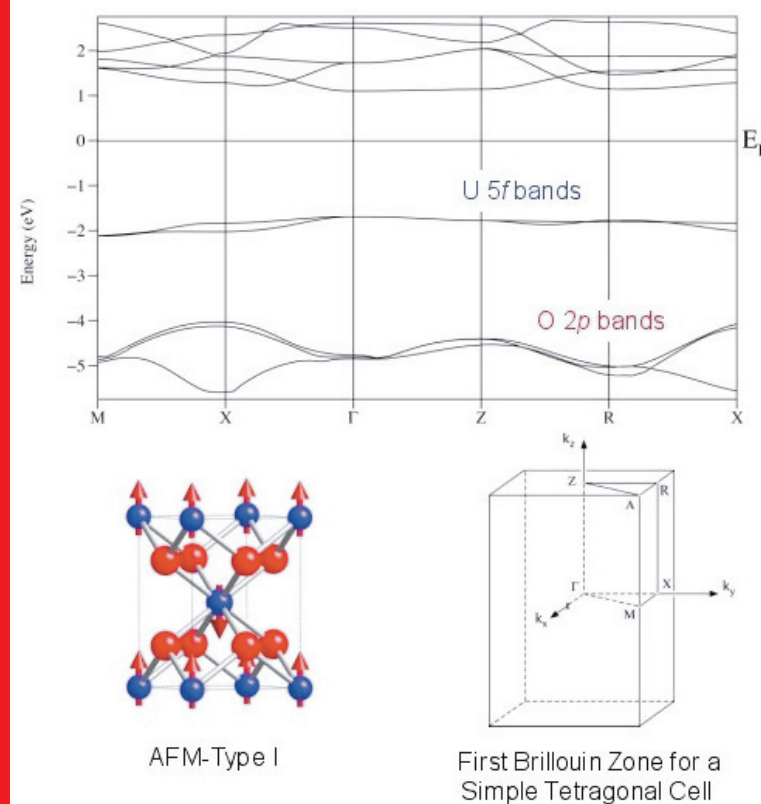


Fig. 2. Alpha-spin band structure of the antiferromagnetic  $\text{UO}_2$  structure (AFM-Type I). The antiferromagnetic cell used in the calculation is a subgroup of the face-centered cubic  $\text{UO}_2$  structure. The red arrows represent the direction of the  $5f^2$  electrons on each uranium atom. The first Brillouin zone with the proper indexing of the relevant points for a simple tetragonal cell is also shown. One can see that the oxygen p bands are well separated from the uranium f bands and that the occupied 5f bands show some dispersion and have a bandwidth of  $\sim 300$  meV.

# Spin-orbit and Jahn-Teller Effects in the Excited States of Platinum II Porphyrins


Cristian V. Diaconu, Richard L. Martin, Enrique R. Batista, T-12; Darryl L. Smith, T-11; Brian K. Crone, Ian H. Campbell, MPA-11; Scott A. Crooker, MPA-NHMFL

The manipulation of the spin degrees of freedom in materials poses fundamental questions for research and represents a very promising area for future device applications. Most work to this point has centered on controlling spin in a traditional electronic semiconductor context. However, organic semiconductors are of great interest as the active layer in spintronic devices because they are composed of light elements with weak spin-orbit interactions and thus are expected to have long spin lifetimes.

In order to understand and manipulate electron spin in organic molecules and devices, we have been developing optical probes that use established spectroscopic techniques for detection. This development entails adding heavy elements with strong spin-orbit coupling to achieve the coupling of spin and optical properties. Thus we also need to understand the optical properties of metal-organic molecules.

The experimental component of this team surveyed a variety of polymer hosts and metal-organic compounds and found two promising materials: metal-organic Platinum(II) Octaethylporphine (PtOEP) (Fig. 1) and the polymer poly (9,9-dioctylfluorenyl-2,7-diyl) (PFO). They then measured the photoluminescence (PL) of PtOEP in PFO matrix as a function of magnetic field and temperature, and found that the zero phonon peak at 643 nm goes away as the temperature is lowered from 20 K to 2 K (Fig. 2). The application of a magnetic field restores the zero phonon peak at 643 nm. Light emission was observed with a net circular polarization from PtOEP and is a clear indication of spin-polarized states in this molecule. The data also exhibits Zeeman splitting and thermal/magnetic activated behavior.


In order to explain the experimental observation we developed a theoretical model. The photoluminescence spectrum of the platinum (II) porphyrins arise from the transition between the lowest triplet excited state and the singlet ground state. The lowest triplet state is spatially doubly degenerate and suffers from a Jahn-Teller instability. We treat the Jahn-Teller problem by reducing the original 105-dimensional Jahn-Teller Hamiltonian to an effective two-dimensional one, whose



effective potential energy surface (PES) determined by hybrid DFT calculations (Fig. 3). By explicit calculation of the intrinsic reaction path connecting the stationary points, we developed an effective one-dimensional PES for the pseudo-rotational motion around the conical intersection (on the bottom of the trough).

From the optimized geometries and associated vibrational frequencies of the singlet ground state, and the excited triplet state, we compute emission transition intensities using the Huang-Rhys approximation. The results of

these calculations, compared with the experimental data, are shown in Fig. 4. The calculation is not broadened by disorder or thermal effects that are implicit in the experimental data. The calculation describes the spectrum well both in the location and relative strengths of the peaks.






Fig. 3. Effective two-dimensional Jahn-Teller PES.




Fig 4. Comparison of experimental and computed photoluminescence spectra of platinum porphine. Experiment: at 85 K, in octane; computation: gas phase, hybrid DFT.

In order to account for the spin-orbit interaction, we developed spin-orbit time-dependent DFT (SO-TD-DFT), which we used to compute the spin-orbit splitting of the first triplet excited state along the effective one-dimensional Jahn-Teller PES. The lowest state is dark everywhere along that path (no singlet mixing with it), and the next spin-orbit state is bright. The splitting between the lowest (dark) spin-orbit state and the bright spin-orbit state varies along the path from 1.3 K to 7.4 K, in agreement with the experimentally observed behavior (Fig. 2). The SO-TD-DFT code will be made available to the community through the Gaussian software.

For further information contact Cristian V. Diaconu at [cvdiaconu@lanl.gov](mailto:cvdiaconu@lanl.gov).

#### Funding Acknowledgments

- Department of Energy, Office of Science, Office of Basic Energy Sciences

# Chronic Beryllium Disease: Structural, Spectroscopic, and Reactive Properties of Aqueous Beryllium Ion

S. Gnanakaran, T-10; Brian L. Scott, MPA-MC

**B**eryllium (Be) is used in industry in metal, oxide, and alloy forms for manufacturing nuclear components, electronic devices, golf clubs, aircraft brakes, x-ray tubes, high temperature ceramics, and structural components in satellites and space shuttles. It also is listed as a Class A environmental protection agency (EPA) carcinogen and causes chronic beryllium disease (CBD), a human granulomatous lung disease. The molecular aspects of beryllium chemistry in the context of biological systems are not well understood. The ability to accurately calculate  $\text{Be}^{2+}$  coordination environments using molecular dynamics (MD) is a crucial step in being able to predict beryllium binding in proteins of interest.

Due to its small size and divalent character, the  $\text{Be}^{2+}$  ion presents a difficult challenge for representing its interactions, classically in an aqueous solution. The charge transfer from water molecules to  $\text{Be}^{2+}$  in this complex, and the interaction with water, may exhibit substantial covalent character. In the past, three-body potentials have been shown to be necessary for obtaining the correct coordination numbers of aqueous  $\text{Be}^{2+}$ . In this vein, we have recently developed a new, simple two-body potential for  $\text{Be}^{2+}$  that has ab-initio-derived Lennard-Jones parameters and a +2 charge.

This new two-body potential captures several structural, reactive, and spectroscopic properties of the  $\text{Be}^{2+}$  ion complex in water. It reproduces the correct radial distribution function and coordination numbers for this cation when compared with published diffraction and nuclear magnetic resonance (NMR) measurements. It also produces a well-established hydrogen bonding between the first and second solvation shells and yields a tetrahedral four-coordinate  $[\text{Be}(\text{H}_2\text{O})_4]^{2+}$  species. When both molecular dynamics and ab-initio results are analyzed together, the role of the solvent in the reactivity of the  $[\text{Be}(\text{H}_2\text{O})_4]^{2+}$  dication is clarified. Namely, the solvent acts to accept protons from the beryllium dication at higher pH values (>3.6). The high reactivity of the resulting  $[\text{Be}(\text{H}_2\text{O})_3(\text{OH})]^+$  species then drives the formation of beryllium hydroxo clusters, such as  $[\text{Be}_2(\text{OH})_3]^+$ ,  $[\text{Be}_3(\text{OH})_3]^{3+}$ ,

$[\text{Be}_5(\text{OH})_6]^{4+}$  and  $[\text{Be}_6(\text{OH})_8]^{2+}$ . Moreover, an ab-initio calculation on a snapshot of the MD calculation gives the best vibration mode energies of aqueous beryllium to date.




Currently, comparative MD studies are being carried out to elucidate the differences between disease and non-disease alleles that could shed much needed microscopic insight on the roles of conformational flexibility and water on binding to  $\text{Be}^{2+}$ . Preliminary results reveal that the effect due to K69E mutation in the disease allele is non-local and cooperative. We also identify potential binding sites and interactions of carboxyl triads that can be exploited for allele-specific genotyping. Details of force field development can be found in [1].

**For further information contact S. Gnanakaran at [gnana@lanl.gov](mailto:gnana@lanl.gov).**

[1] S. Gnanakaran et al., *J. Phys. Chem. B*, in press (2008).


## Funding Acknowledgments

- Los Alamos National Laboratory Directed Research and Development Program



**Fig. 1. Top:** Structural characteristics of the hydrated beryllium complex,  $[\text{Be}(\text{H}_2\text{O})_4]^{2+}$ . The beryllium ion-water oxygen radial distribution function (solid) and the corresponding integration numbers (dashed) are shown. **Insert:** Geometric values of  $[\text{Be}(\text{H}_2\text{O})_4]^{2+}$  complex. Bond and angular values are marked. **Bottom-left:** Interplay between first and second solvent shells. Hydrogen bonding network between first and second shell water molecules. The Be cation is in red. The color gradient from green to blue is used to mark the distance from the cation (green:close and blue:far). **Bottom-right:** Influence of second solvent shell on the geometry and charge distribution of the  $[\text{Be}(\text{H}_2\text{O})_4]^{2+}$  complex. The figure provides results from *ab-initio* calculations. The numbers in the brackets are O-O distances and the numbers in parenthesis correspond to distances from a configuration obtained from MD.





# Climate Modeling

It is hard to imagine a discipline with more impact within the computational sciences than that of climate modeling. LANL has a world-class program for global climate and ocean modeling, built over many years. As part of the global climate simulation community, LANL scientists and their models had a significant impact in the Intergovernmental Panel on Climate Change (IPCC), which was awarded the Nobel Peace Prize in December 2007. LANL is fully responsible for all advanced numerical models for ocean and sea ice that are used in global climate change predictions. Such simulations play an important role in a variety of policy issues that affect our nation and the world. The coupling of the ocean to the climate is critical to a holistic understanding of climate and, more importantly, how the climate is changing and what the changes mean.

The papers in this section describe the Climate, Ocean, and Sea Ice Modeling Project at LANL and its focus and impact. Ongoing efforts highlight new discoveries that allow much higher resolution simulations of the ocean, new methods to help understand eddies within the ocean and their effect, and new understanding about abrupt climate change. Researchers are looking for new indicators of change from melting ice effects in high latitudes and simulating the biogeochemistry of the ocean. This is truly world-class science that has a positive impact on the global community.

# The LANS-alpha Turbulence Parameterization in Ocean-climate Modeling

Matthew W. Hecht, Mark R. Petersen, Beth A. Wingate, CCS-2;  
Darryl D. Holm, CCS-2, Imperial College, London

Ocean-climate models are typically run at low resolution (1 deg., or 100 km grid cells) in climate simulations due to the computational requirements of the coupled components and duration of the simulations. This resolution is well above the Rossby radius of deformation over most of the ocean, the typical horizontal size of eddies in the ocean. As a result, ocean-climate simulations only include the mean, large-scale flow, but not the eddies that occur in ocean observations (Fig. 1). These eddies affect the mean circulation by transporting heat, salinity, and kinetic energy.

The goal of turbulence modeling is to capture the effects of small scale structures on the large-scale flow. The Lagrangian-averaged Navier-Stokes alpha (LANS-alpha) model, developed in part by Darryl Holm, is a turbulence parameterization that has been shown to improve turbulence statistics at low resolution. It has previously been tested in turbulent pipe flow, large eddy simulations, shallow water models, and quasigeostrophic models, all with positive results.

We have implemented the LANS-alpha model in the Parallel Ocean Program (POP) [1], a code developed at LANL for climate change simulations. Several new technical challenges had to be overcome in this work, including designing an efficient algorithm that fits into POP's split time-stepping scheme [2], and choosing an efficient and stable smoothing method for the advective velocity [3].

LANS-alpha improves the representation of turbulence in all statistics measured in an idealized channel configuration [4]. Figure 2 shows that POP with LANS-alpha (POP-alpha) produces eddy kinetic energy and vertical temperature profiles that are similar to doubled-resolution simulations without LANS-alpha. This implies a great savings in computational time, as doubling the horizontal resolution requires ten times longer, while adding the LANS-alpha model only takes 30% longer. This should have a large impact in climate change simulations, as the ocean is currently forced to run at low resolution due to the demanding computational requirements of climate models.

For further information contact Mark R. Petersen at [mpetersen@lanl.gov](mailto:mpetersen@lanl.gov).

- [1] D.D. Holm et al., *Los Alamos Science* **29**, 152-172 (2005).
- [2] M.W. Hecht et al., submitted to *J. Comp. Phys.* (2008).
- [3] M.R. Petersen et al., *J. Comp. Phys.* (2008), in press.
- [4] M.W. Hecht et al., submitted to *J. Phys. A.* (2008).

## Funding Acknowledgments

- Department of Energy, Office of Biological and Environmental Research, Climate Change Prediction Program
- Los Alamos National Laboratory Center for Nonlinear Studies
- Los Alamos National Laboratory Directed Research and Development Program

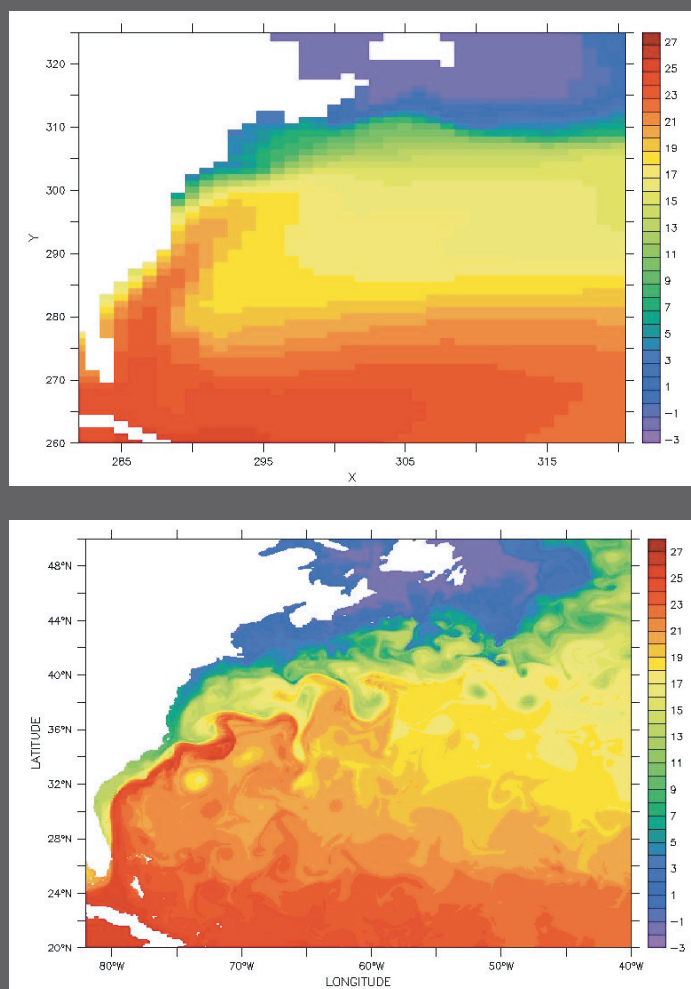


Fig. 1. Surface temperature (in deg. C) of an ocean simulation using a horizontal grid resolution used in climate modeling (1.0 deg., top) and a high resolution ocean simulation (0.1 deg., bottom). Climate model simulations do not resolve the Rossby Radius, and so do not include eddies and their associated turbulent fluxes.

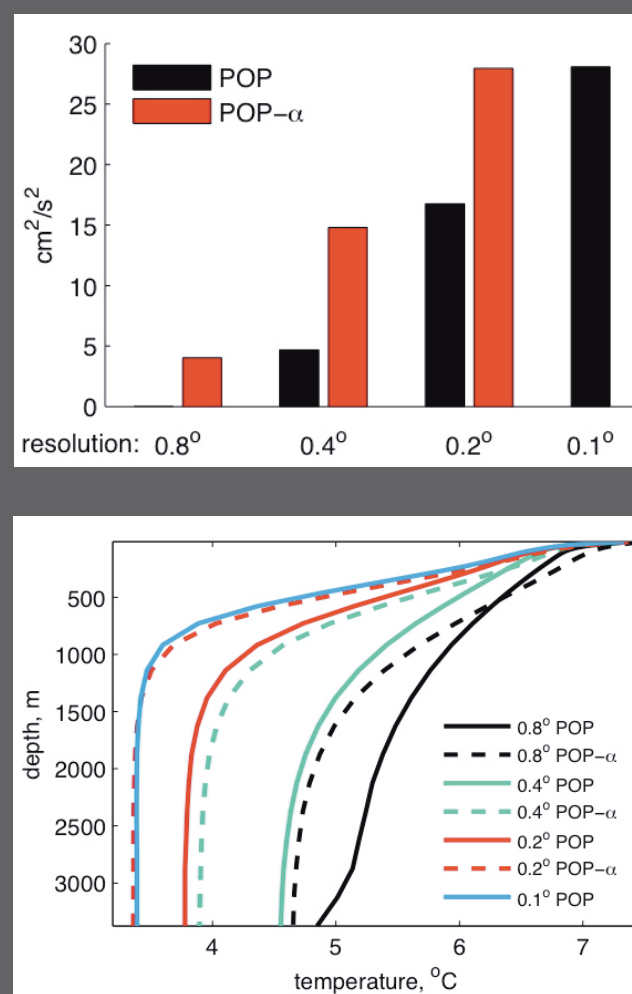


Fig. 2. Comparison of turbulence statistics produced by POP and POP with LANS-alpha (POP-alpha). With LANS-alpha, both eddy kinetic energy (top) and vertical temperature profiles (bottom) resemble doubled-resolution simulations without LANS-alpha. LANS-alpha is an effective turbulence parameterization because it is much cheaper computationally than doubling the resolution.

# A Review of North Atlantic Ocean Modeling in an Eddying Regime

Matthew W. Hecht, CCS-2

The Climate, Ocean, and Sea Ice Modeling project (COSIM) [1] began as an effort to bring ocean general circulation modeling to massively parallel computers, and quickly expanded to include applications. The second large application undertaken by the project was a high-resolution simulation of the North Atlantic Ocean. Initiated in 1996 (before the advent of the moniker COSIM), the effort led to publication of a paper [2], which is now seen as marking the advent of the era of ocean modeling in a strongly eddying regime.

Eddies within the ocean are the analog to storm systems in the atmosphere. Their sizes are determined by wave propagation speeds and the rotational period of the Earth. The slower speed of the relevant waves in the ocean makes for a much smaller size than that found in the atmosphere, and tremendously high spatial resolution is therefore required in order to adequately resolve oceanic eddies. The importance of the eddies lies not only in the mixing of water properties, but in their feedback on the large-scale, time-mean circulation of the oceans.

A decade ago the weak level and inadequate influence of eddy variability on mean circulation was most prominently illustrated by the poor representation of the North Atlantic's Gulf Stream. The simulation presented in [2], performed on the Connection Machine, CM-5, in what was then known as LANL's Advanced Computing Laboratory, not only achieved a level of eddy variability that compared favorably with that inferred from satellite-borne instruments, but also produced a greatly improved representation of the Gulf Stream (Fig. 1).

North Atlantic Ocean modeling has continued to be an active focus for the Laboratory's COSIM project, with research focused both on physical understanding and on the more empirical problem of how to best configure a model for the "grand challenge-scale" problem of global ocean modeling in a strongly eddying regime.

The progress of this past decade is discussed in a review paper [3], to appear early in 2008. One point emphasized in the review is the importance of eddies in invigorating the deep circulation (Fig. 2). This deep circulation in turn interacts with the more readily observed currents of the upper ocean.

The North Atlantic Ocean is also of central interest to the Earth's climate. New work is underway to determine the importance of a more accurate and realistic circulation to the stability of modeled climate.

**For further information contact Matthew Hecht at [mhecht@lanl.gov](mailto:mhecht@lanl.gov).**

[1] <http://climate.lanl.gov/COSIMHistory.pdf>

[2] R. Smith et al., *J. of Physical Oceanography*, **30**, 1532-1561 (2000).

[3] M. Hecht and R. Smith, in *Ocean Modeling in an Eddying Regime*, American Geophysical Union Geophysical Monograph Series, Hecht and Hasumi, Eds. (2008).

## Funding Acknowledgments

- Department of Energy, Office of Biological and Environmental Research, Climate Change Prediction Program



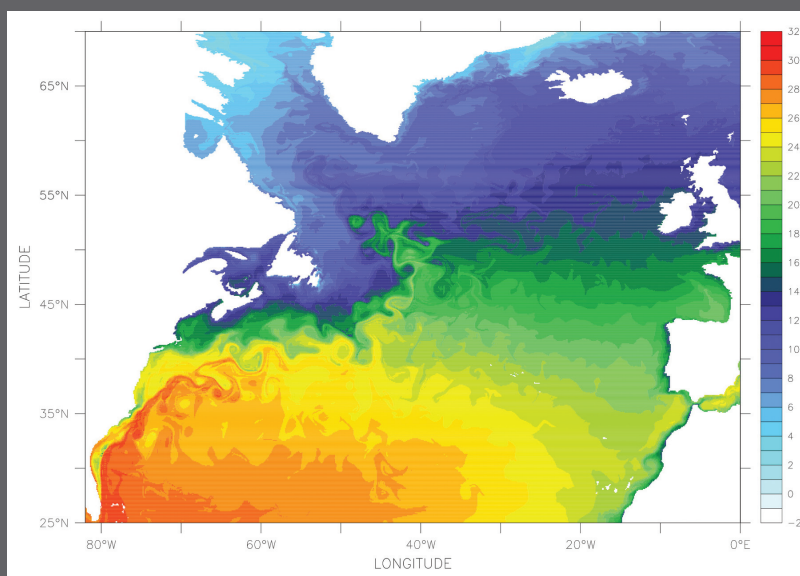


Fig. 1. Sea surface temperature in the Gulf Stream region of the North Atlantic Ocean, from a simulation at the relatively high horizontal grid resolution of 0.1 degree.




Fig. 2. Deep ocean transport from the same simulation as in Fig. 1. Total volume transport of cold, deep waters with potential temperatures between 3 and 4 degrees Celsius is shown, with transport vectors superimposed. This vigorous deep circulation is now known to be one of the elements required in order to establish the correct upper ocean circulation of Fig. 1.

# Advanced Numerical Techniques and Algorithms to Study Abrupt Climate Change

Balu Nadiga, CCS-2

Past records show that abrupt climate changes were most common when the climate system was being forced to change most rapidly or when the thermohaline circulation (THC) was weak [1]. So, the question is whether the rapid climate change that we are presently witnessing (Fig. 1), along with model predictions of a weakening of the THC [2], increases the chances of abrupt climate change in the near future.

Abrupt climate change is a manifestation of complex nonlinear chaotic behavior in the climate system and occurs when the climate system is forced to cross a threshold, leading to a transition to a new climate state at a rate that is faster than the cause and which is determined internally by the system (e.g., see Fig. 2). Progress was made in understanding the dynamics underlying such complex nonlinear chaotic behavior by using the simplest settings of low-dimensional differential equations and maps, mostly in the 70s and 80s. An extension of techniques that were developed and used for low-dimensional systems (a popular example is the package AUTO) to high-dimensional systems has had to await more recent advances in computational techniques and is only now becoming feasible. To wit, the philosophy and




Fig. 1. The rise in global temperature is one aspect of the unprecedented rapid change in climate that we are presently experiencing. From [1].

methodology underlying present day oceanic general circulation models (OGCMs), most of which had their origins in the 70s, make them more suited for studying gradual rather than abrupt changes.

In order to help policy makers make informed decisions about safe levels of greenhouse gases in the atmosphere, climate scientists need to provide probabilistic estimates that the climate system will cross some threshold leading to abrupt climate change and probabilistic estimates of the consequences of the resulting climate shift on various aspects of the Earth system. Factors contributing to the uncertainty associated with abrupt climate change involve 1) uncertainty about the location of thresholds as a function of uncertain model parameters and parameterizations, 2) uncertainty about the location of the present climate system in parameter space, and 3) uncertainty about future climate forcing including, for example, future greenhouse gas emissions. Quantifying the risks associated with abrupt climate change therefore requires the systematic exploration of the positions of various nonlinear thresholds as a function of uncertain model parameters and forcing.

Techniques required to study abrupt changes are complementary to those used in present day ocean models. For example, present day OGCMs cannot compute or track unstable equilibria—the simplest dynamical objects that underlie abrupt changes. So also, thermodynamic spinup of present day OGCMs is extremely slow: The time step which is determined by the fastest gravity wave speeds is between a few seconds to a few hours depending on spatial resolution, whereas the times associated with setting up deep circulation is of the order of thousands of years.




Fig. 2. Two regimes of Stommel's box model of THC. When diffusion is weak (orange), the model exhibits abrupt change and hysteresis, but not when diffusion is strong (green). From [1].




Fig. 3. Modeled sea surface height in the North Atlantic. With the new techniques, a time step of a day was used. The time step using the traditional methodology is limited to about an hour.

height as modeled in a simplified setup of the North Atlantic. In this case, the maximum timestep in the traditional methodology is an hour, whereas the present computation used a timestep of a day. Other advantages include a consistent and uniform treatment of terms in the governing equations. Work is underway to further improve the efficiency of these schemes. As an example, a high degree of improvement in efficiency is achieved (Fig. 4) by “preconditioning” the gravity waves, which happen to be the fastest waves in the system. Thus with this and other related techniques that we are presently working on, we expect to significantly ameliorate the spinup problem as well.

**For more information contact Balu Nadiga at [balu@lanl.gov](mailto:balu@lanl.gov).**

Collaborators: Mark Taylor, Sandia National Laboratories;  
Francois Primeau, Univ. of California, Irvine;  
Jens Lorenz, Univ. of New Mexico

We have now developed new and complementary numerics based on Jacobian Free Newton Krylov techniques in a popular OGCM that enables studying thresholds and abrupt change scenarios. That is, with this approach, we can now track changes in ocean circulation as key parameters are changed. This approach has been implemented in various other problems of interest to the Laboratory as in [4].

Further, when used in place of traditional time stepping algorithms, this method allows for a time step that is of relevance to the physical phenomenon that is being studied. The time step is not limited by the fastest modes of the system [3]. For example, Fig. 3 shows the sea surface




Fig. 4. Impact of a preconditioner on the efficiency of the Newton-Krylov methodology. The upper line shows the number of iterations of generalized minimal residual method (GMRES) required without the preconditioner and the lower line with the preconditioner. GMRES is a costly component of the scheme.

- [1] Abrupt Climate Change: Inevitable Surprises, National Research Council, (2002).
- [2] Intergovernmental Panel on Climate Change, Fourth Assessment Report, (2007); <http://www.ipcc.ch>
- [3] B.T. Nadiga et al., Computer and Mathematical Modeling, 44, 870-886, (2006).
- [4] D.A. Knoll et al., Journal of Scientific Computing, 25, 213-30, (2005).

#### Funding Acknowledgments

- Department of Energy, Office of Biological and Environmental Research
- Department of Energy, Office of Biological and Environmental Research, Climate Change Prediction Program
- Los Alamos National Laboratory Directed Research and Development Program

# Ocean Modeling and the Representation of Unresolved Scales

Balu Nadiga, CCS-2

Ocean circulation is a prototypical “Strongly Coupled Multiscale Phenomenon” wherein a vast range of spatial and temporal scales interact to produce intricate and complex behavior that drives low-frequency climate variability. The computational challenge of modeling ocean circulation is to accurately represent this vast range of spatial and temporal scales and their interactions.

State-of-the-art climate studies use realistic ocean models but at coarse resolutions. Comparisons of modeled circulation at these resolutions with satellite observations and in-situ measurements show that the modeled circulation and its variability are both pathologically sluggish.

The separation of the Gulf Stream from the East coast of the United States in models highlights the issue. In Fig. 1 [1], while at the higher resolution on the left, the Gulf Stream separates correctly at Cape Hatteras and the simulation correctly captures dynamics of the separated jet further downstream; most of these features are incorrect in the simulation on the right. The latter simulation uses reduced resolution, which is still higher than resolutions used in climate studies. The importance of the unresolved small-scales is highlighted by the fact that the reduced-resolution simulation is configured identically to the higher-resolution simulation.

A new approach to modeling unresolved subgrid-scales [2] in simulations such as these is based on the idea of regularizing small resolved scales. This is done by modifying the nonlinear cascade processes so that the importance of small scales is deemphasized [3,4,5]. An attendant feature of this regularization is a modification of the dispersion relation for short waves in the system [5,6] that ameliorates time step constraints [6]. This approach is in contrast to the traditional approach of modeling the effects of unresolved scales on resolved scales as a dissipative downgradient closure. The chief advantage of

the new approach over the traditional downgradient closure lies in the ability of the new approach to represent backscatter [7]—a nondissipative influence of the small unresolved scales on the resolved scales—besides the usual damping effect of the small unresolved scales.

Figure 2 shows a simplified setup that highlights the applicability of this approach to the problem of Gulf Stream separation. Again, the higher resolution simulation on the left shows proper separation at Cape Hatteras, but the reduced-resolution run (center) displays the characteristic stationary anticyclonic meander that prevents proper separation. On the other hand, at the reduced resolution, with the new regularization-based subgrid model on the right, the stationary anticyclonic meander is absent and the overall representation of the dynamics of the Gulf Stream separation is improved.

This improvement is due to an enhancement of the inverse-cascade of energy, a characteristic feature of large-scale flows. The enhanced inverse-cascade makes the large-scale flow more inviscid, as shown in an idealized setting in Fig. 3. In that figure, the distribution of energy with scale is shown on a log-log plot for different values of the regularization parameter  $\alpha$ . All the flows are fully resolved. The solid line corresponds to the original (un-regularized) system. Dotted, dashed, and dot-dashed lines correspond to increasing values of  $\alpha$ . Note the de-emphasizing of small scales with increasing  $\alpha$ . Inverse energy cascade results when energy flows from the forcing scale ( $k=10$ ) to larger scales. That regularization of small

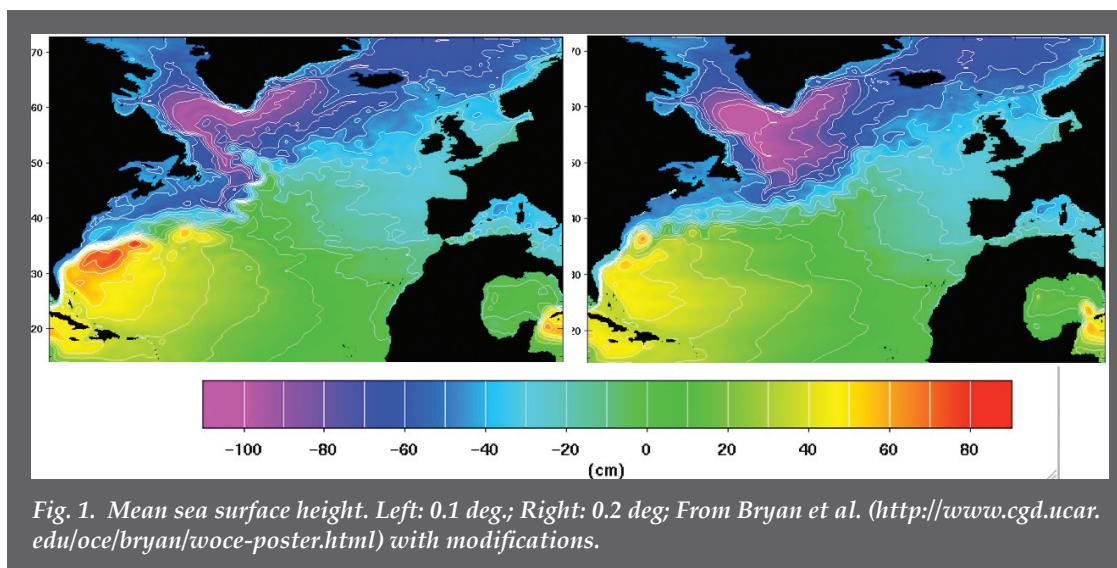



Fig. 1. Mean sea surface height. Left: 0.1 deg.; Right: 0.2 deg; From Bryan et al. (<http://www.cgd.ucar.edu/oce/bryan/woce-poster.html>) with modifications.





scales results in an enhancement of the inverse-cascade of energy to large-scales also demonstrates the important interplay between numerics and physics in the simulation of these systems.


Work is underway to a) further understand and improve the subgrid model, and b) to implement such subgrid models in ocean models used in IPCC-class climate simulations [8].

**For more information contact Balu Nadiga at [balu@lanl.gov](mailto:balu@lanl.gov).**

- [1] R.D. Smith, et al., *Journal of Physical Oceanography*, **30**, 1532-1561 (2000).
- [2] D.D. Holm, J.E. Marsden, and T.S. Ratiu, *Phys. Rev. Lett.* **80**, 4273-4277 (1998).
- [3] B.T. Nadiga and S. Shkoller, *Phys. Fluids*, **13**, 1528-1531 (2001).
- [4] B.T. Nadiga and L.G. Margolin, *J. Phys. Ocean*, **31**, 2525-2531 (2001).
- [5] D.D. Holm and B.T. Nadiga, *J. Phys. Ocean*, **33**, 2355-2365 (2003).
- [6] B.A. Wingate, *Mon. Weather Rev.*, **132**, 2719-2731 (2004).
- [7] B.T. Nadiga, Orientation of eddy fluxes in geostrophic turbulence, submitted to *Proc. Royal Soc. A*, (2007).
- [8] M.W. Hecht, et al., Implementation of the LANS-alpha turbulence model in a primitive equation ocean model, submitted to *J. Comp. Phys.*, (2007).

#### Funding Acknowledgments

- Department of Energy, Office of Biological and Environmental Research, Climate Change Prediction Program
- Los Alamos National Laboratory Directed Research and Development Program





## New Result Shows Separation of Time Scales in Rotating and Stratified Flows at High Latitudes

Beth Wingate, CCS-2; Miranda Holmes, New York Univ.; Mark Taylor, Sandia National Laboratories

Earth's high latitudes (Fig. 1) stand to be among the first regions affected by climate change issues due to changes induced by melting ice in the Arctic and Antarctic. The dynamics in the high latitudes is affected by the higher rotation rate as measured by the Rossby number. The smaller the Rossby number, the more important the rotation is to the dynamics. Motivated by gaining fundamental understanding of ocean dynamics at high latitudes, we have derived new results based on the method of multiple scales presented in [1], which address the scale separation between slow- and fast-time scale dynamics in the limit of fast rotation while retaining order one effects due to stratification. This paper has been submitted for publication [2]. The result is important because it leads to new equations for the slow dynamics.

The principle results are:

- Finding that as the rotation rate increases, the fast and slow dynamics decouple.
- Development of new equations for the slow dynamics that include conservation laws of the slow variables. One important consequence of the new slow equations is that the horizontal and vertical kinetic energy decouple and conserve their kinetic energy independently.

Figure 2 shows the potential enstrophy as a function of the Rossby number for numerical simulations of high wave number forcing numerical experiments [3]. As the Rossby number decreases (or the rotation rate increases), the potential enstrophy becomes entirely slow. Figure 3 shows the horizontal kinetic energy as a function of the Rossby number. As the Rossby number decreases, the time rate of change of the vertical kinetic energy approaches zero.

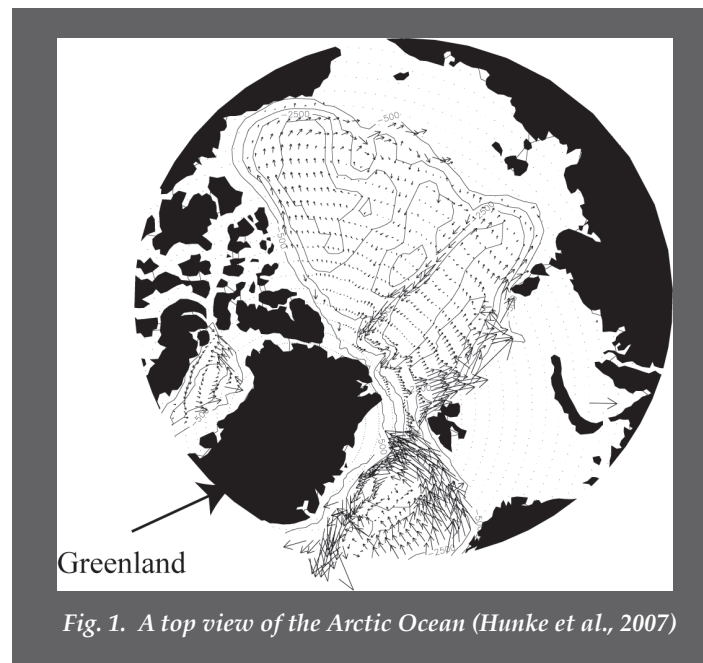


Fig. 1. A top view of the Arctic Ocean (Hunke et al., 2007)

For further information contact Beth Wingate at [wingate@lanl.gov](mailto:wingate@lanl.gov).

- [1] P. Embid and A. Majda, *Geophysical and Astrophysical Fluid Dynamics*, **87**, 1-50 (1998).
- [2] B. Wingate, M. Holmes, M. Taylor submitted to *Journal of Fluid Mechanics* (2007).
- [3] L. Smith and F. Waleffe, *Journal of Fluid Mechanics*, **451**, 145-168 (2002).

#### Reference

- E. Hunke and M. Holland, *Journal of Geophysical Research*, **112**, c06S14 (2007).

#### Funding Acknowledgments

- Department of Energy, Office of Biological and Environmental Research, Climate Change Prediction Program
- Department of Energy, Office of Science




Fig. 2. The potential enstrophy as a function of Rossby number. As the Rossby number decreases, the total potential enstrophy becomes entirely slow.




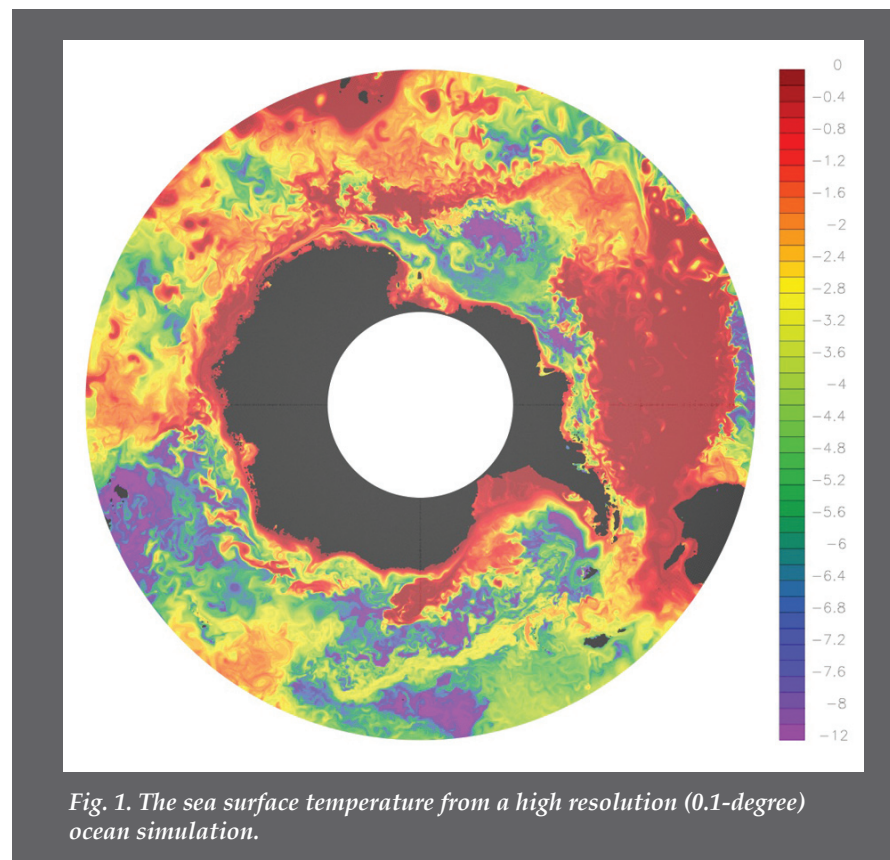
Fig. 3. The time rate of change of the horizontal kinetic energy approaches zero as the Rossby number decreases.

## Climate, Ocean, and Sea Ice Modeling


Philip Jones, Elizabeth Hunke, Bill Lipscomb, Matthew Maltrud, Todd Ringler, Bin Zhao, T-3; Scott Elliott, Matthew W. Hecht, Balu Nadiga, Nicole Jeffery, Mark R. Petersen, Wilbert Weijer, Beth A. Wingate, CCS-2

The Climate, Ocean, and Sea Ice Modeling (COSIM) project develops and maintains advanced numerical models of the ocean and sea ice for use in global climate change projections. COSIM models were used extensively in simulations for the recent climate assessment assembled by the Intergovernmental Panel on Climate Change (IPCC) that was awarded the 2007 Nobel Peace Prize. The flagship Parallel Ocean Program (POP) and Community Ice CodE (CICE) are the ocean and sea ice components of the Community Climate System Model that contributed the largest number of simulations to the recent IPCC assessment, and are used by a large number of climate scientists worldwide. In addition to ocean and sea ice models, COSIM scientists are working on new ice sheet models to more accurately simulate and predict the fate of large ice sheets covering Greenland and West Antarctica. Ocean ecosystems and their impacts on the carbon and sulfur cycles have been another important addition to the physical models.

While the emphasis is on developing ocean and ice models for the climate community, COSIM researchers also use the models to address scientific issues. Of particular interest over the past several years has been the impact of ocean eddies on the global ocean circulation. Ocean mesoscale eddies are about 20 km in size, and resolving these eddies requires the highest-resolution ocean simulations and high-performance computing capabilities. COSIM ocean simulations have demonstrated dramatic improvement in the representation of ocean circulation, including realistic representations of the North Atlantic current system and Gulf Stream separation. Past high-resolution simulations have focused on the ocean, but COSIM team members are in the process of performing such simulations in a fully coupled climate model with eddy-resolving ocean and sea ice components joined together with high-resolution atmosphere and land models.



Over the next few years, polar climate change and its global impacts will be an important emphasis as polar regions experience the largest response to climate change. Evaluating the possibility of ice sheet collapse, simulating the observed rapid Arctic sea ice melt, expanding the representation of polar ecosystems to include ice organisms, identifying thresholds of ocean circulation stability resulting from large freshwater inputs, and quantifying the magnitude and rate of sea level rise will all be important areas of research as we assess the impacts of climate change.



*Fig. 2. Concentration of a passive dye tracer at 317 m depth after 7 months of simulation. This tracer has arbitrary units and its concentration is continuously reset to be 1 in the top level.*

For further information contact Philip Jones at [pwjones@lanl.gov](mailto:pwjones@lanl.gov).

#### **Funding Acknowledgments**

- Department of Energy, Office of Biological and Environmental Research, Climate Change Prediction Program
- Department of Energy, Office of Science, Scientific Discovery through Advanced Computing Program

## Biogeochemistry in Los Alamos Climate, Ocean, and Sea Ice Modeling


Scott Elliott, CCS-2; Mathew Maltrud, Elizabeth Hunke, T-3

Geocycling of carbon by the marine biota determines to a large extent the state of planetary-scale climate and, further, its response to anthropogenic greenhouse gas release. The Climate, Ocean and Sea Ice Modeling (COSIM) project has responded to the challenge of comprehending this situation by developing and running global ocean biogeochemistry simulations. This process is accomplished by adding several dozen new tracers to the Los Alamos Parallel Ocean Program (POP). Nutrients, phytoplankton, zooplankton, bacteria, and organic detritus are all represented along with their source sink interrelationships. In essence, the COSIM team models the entire marine food web and its integrated feedbacks on climate. Geocycling processes considered include: a) large scale ecodynamics and the associated carbon drawdown, b) limitations placed upon primary production by trace forms of iron or nitrogen, and c) the role of oceanic organisms in controlling aerosol/cloud reflectivities via gas exchange.

Most biogeochemistry research conducted thus far in POP has focused on computation of phytoplankton distributions with links to carbon and iron because they are the elements most directly related to climate management [1,2]. The sulfur cycle feeds back on carbon dioxide-driven warming through sea-air transfer of the trace gas dimethyl sulfide (DMS). COSIM has developed detailed DMS processing mechanisms which have fared well in recent intercomparisons [3]. A host of other volatile species has been modeled in parallel. Examples include carbon monoxide as a critical surface ocean photochemical, carbonyl sulfide, nitrous oxide, ammonia, and a variety of pure and halogenated hydrocarbons [4].

With wholesale Arctic environmental change now at hand, COSIM biogeochemists are embarking on an effort to develop polar and general high-latitude marine systems models. For example, ice algae have now been added to the global ecodynamics structure. Interactions with sea floor sediment, the possibility of methane clathrate destabilization there, and the unique organisms of cold surface

waters will all be considered. An explosion of data is now becoming available to marine systems modelers in the form of metagenomes. COSIM biogeochemists are collaborating with theoretical biologists and genomics specialists at the Laboratory in order to develop methods for informing large-scale models of the oceanic metabolism.









Fig. 2. Ice algal patterns simulated for late spring of 1982, in the Los Alamos Community Ice CodeE (CICE) package, which can be coupled to POP. The color bar is linear in concentration and runs from 0 (deep blue) through 10 (red) in the units millimoles of Si per square meter – the organisms are a form of pennate diatom constructing silicate frustules (shells).

For further information contact Scott Elliott at [sme@lanl.gov](mailto:sme@lanl.gov).

- [1] S. Chu, S. Elliott, and M. Maltrud, *Chemosphere*, **50**, 223-235 (2003).
- [2] L. M. Rothstein et al., *Oceanography*, **19**, 2-51 (2006).
- [3] S. Elliott, S. Chu, and D. J. Erickson, *Environmental Modeling and Software*, **22**, 349-358 (2007).
- [4] S. Elliott, S. Chu, and C. Dean, TRACEGAS\_MOD: Geochemical Processing for Low Concentration Volatiles in the CCSM Ocean in *Environmental Sciences and Environmental Computing*, Vol. 3, FiatLux Publications, San Francisco (2008).

#### Funding Acknowledgments

- Department of Energy, Office of Biological and Environmental Research, Climate Change Prediction Program
- Department of Energy, Office of Science, Scientific Discovery through Advanced Computing Program





# Information Science and Technology

We are witnessing the emergence of powerful tools for dealing with both enormous amounts of data and near-instantaneous access to data from disparate sources. A new discipline is emerging, Information Science and Technology, that draws upon computer science and mathematics and extends these disciplines to enable us to understand and obtain insight from large and disparate sources of data. LANL researchers are helping to shape this new discipline as a powerful tool for science.

This section includes contributions of Information Science and Technology to a wide range of application areas ranging from increasing the effectiveness of radiation sensors employed at transportation portals to understanding human brain activity. These articles describe approaches for developing understanding from large volumes of sensor data and combining data from different sources to obtain new insights. Also included are articles on expanding the science and technology that underpin the rapidly evolving field of Information Science and Technology. The last paper in this section describes a broad-based collaboration between National Laboratories and Universities and addresses some of the technical challenges of Information Science and Technology.

# RaveGrid: Raster to Vector Graphics for Image Data

Sriram Swaminarayan, CCS-2; Lakshman Prasad, ISR-2

We present RaveGrid, a software that efficiently converts a raster image to a scalable vector image comprising polygons whose boundaries conform to the edges in the image. The resulting vector image has good visual quality and fidelity and can be displayed at various sizes and on various display screen resolutions. The software can render vector images in scalable vector graphics (SVG) or encapsulated postscript (EPS) formats. The ubiquity of images, on the Web and in communications, as well as the wide range of devices, from big-screen televisions to cellular phones that support image display, calls for a scalable and more manipulable representation of imagery. Moreover, with the growing need for automating image-based search, object recognition, and image understanding, it is desirable to represent image content at a semantically higher level by means of tokens that support computer vision tasks.

**Algorithm.** Our software is based on *VISTA*, a broad algorithmic framework [1] for performing image segmentation that exploits both region and edge cues inherent in images. Perceptual organization of regions and edges into polygonal segments is modeled using the proximity-based regional relationships between edges established by Delaunay triangulations. More precisely, a digital image is processed to extract edge pixel chains and a constrained Delaunay triangulation of the edge contour set is performed to yield triangles that tile the image without crossing edge contours. Each triangle is attributed a color by sampling pixels within it. A combination of rules, each of which models an elementary perceptual grouping criterion such as proximity or continuity, determines which adjacent triangles should be merged. A grouping graph is formed with vertices that represent triangles and edges between vertices, and that correspond to adjacent triangles to be merged according to the combination of grouping rules. Connected components of the grouping graph then yield collections of triangles that form polygons that segment and vectorize the image.

**Implementation.** Our software takes as input a digital raster image comprising pixels and an optional quality parameter—the level of detail (LOD) that controls the hysteresis thresholds in a Canny edge detector. The steps of the algorithm

sketched above are efficiently implemented in C, and a graphical user interface (GUI) displays both the raster input image and the vector output image. The details of the vectorization can be compared by synchronous zoom and pan functionalities. The result may be stored as a gzip compressed SVG or EPS document.

**Performance.** RaveGrid currently vectorizes an average of two VGA resolution (640x480) frames per second (or about 1.75 s/MP) on a 2.5 GHz Pentium 4 CPU with 1 GB of RAM, scaling linearly to over 30 megapixels (Fig. 1). At the lowest quality setting, it serves as an image segmenter, while at higher quality settings it renders high visual quality vector images. Figure 2 shows a sample image of red and green chiles vectorized at different LODs: high, medium, and low. At the high LOD the vectorized image is indistinguishable from the original. As we lower the LOD, even though fine details and highlights disappear, the perceptual content of the image is preserved making this approach amenable to multi-scale object analysis.


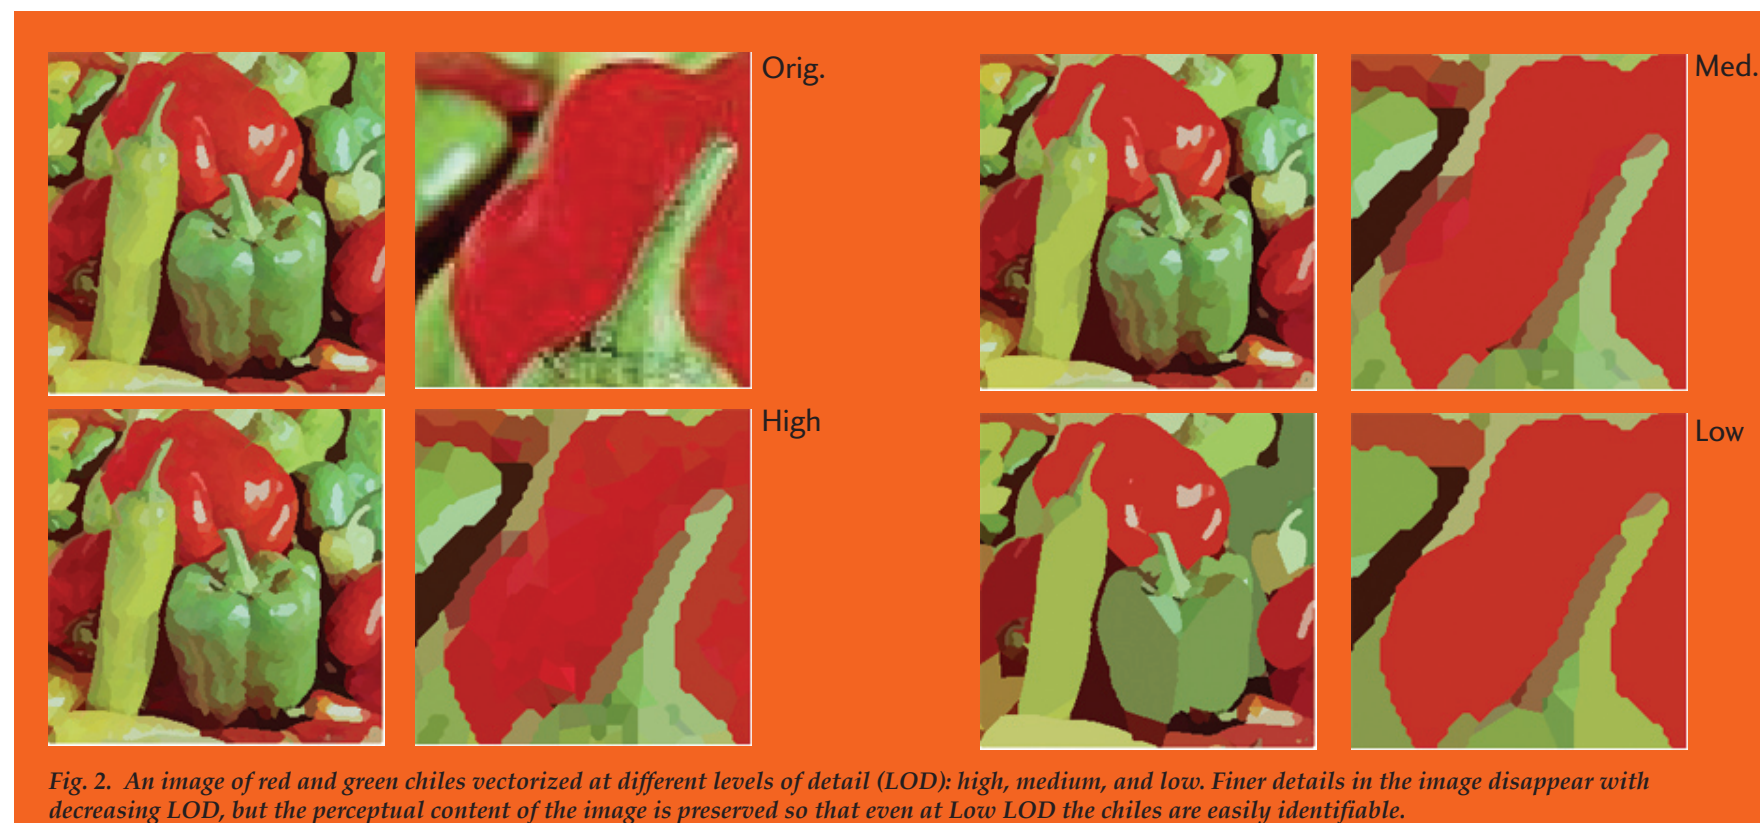


Fig. 1. Scaling of RaveGrid as a function of image size.





For further information contact Sriram Swaminarayan at [sriram@lanl.gov](mailto:sriram@lanl.gov).

[1] L. Prasad and A. Skourikhine, *Pattern Recognition*, **39**, 501-514 (2006).

#### Funding Acknowledgments

- Department of Energy, National Nuclear Security Administration, Advanced Simulation and Computing Program



# Effective and Efficient Visualization and Analysis of Large-scale Data on Emerging Computer Architectures

Patrick S. McCormick, James Ahrens, CCS-1; Jeff Inman, HPC-1

Simulation and modeling play a fundamental role in many missions of the Department of Energy. These applications are producing data at ever-increasing rates that are overwhelming our capabilities to explore, hypothesize, document, and thus fully interpret the underlying details. These tasks will continue to become increasingly difficult as we make advancements towards petascale computing.

The computer architectures that support these computations are undergoing a revolutionary change as manufacturers transition to building chips that use an increasing number of processor cores. In addition, graphics hardware that was once designed entirely for rendering has rapidly evolved into a powerful general-purpose processor. While these trends have the ability to provide new capabilities and increase performance, they will place a significant strain on software development activities. The goal of our effort is to explore techniques that leverage the power provided by these emerging processor architectures to help scientists efficiently and effectively understand large amounts of information produced by computational simulations.

The fundamental challenge is how to provide scientists with a set of software tools that allow them to leverage the power of emerging computer architectures without requiring a detailed understanding of the available processors.

Driven by their wide use within the entertainment industry, graphics processors (GPUs) are among the most cost-effective, high-performance processors available today. Figure 1 provides a summary of the performance trends and prices of GPUs and CPUs as of August 2007. Similar trends are also evident in the design of other processor technologies. For example, IBM's Cell Broadband Engine is capable of 256 billion floating point operations per second [1]. Unfortunately, these processors are complex and can be very difficult to program efficiently.






Fig. 1. Historical GPU versus CPU performance trends and current pricing for the most recent processors. Performance numbers courtesy of Mike Houston and John Owens. Pricing data from <http://www.price-watch.com>, August 2007.

To address this issue we have developed a parallel programming language targeted at supporting visualization, data analysis, and general-purpose computation. The Scout language allows visualization tasks, which are often represented by static user interfaces, to be directly expressed in programmatic form [2]. In addition, the supporting runtime and compiler infrastructure provides support for leveraging the power of multiple forms of parallelism and the potential for supporting computer systems designed using a collection of heterogeneous processors [3]. Figures 2 and 3 show example computations created using Scout.



*Fig. 2. Volume-rendered results of two selected entropy ranges colored by corresponding velocity magnitude. Both the entropy and velocity magnitude were computed directly using Scout. Performance was approximately 12 times faster than a 3.0 GHz Intel Xeon EM64T processor.*



*Fig. 3. A qualitative comparative visualization of the results produced by two different cosmology simulations codes studying the nonlinear structures in the dark and luminous matter that makes up the universe. The results were accelerated by leveraging the power of the graphics processor of a laptop computer. Using this approach we were able to compute, query, and render 32 million particles at 24 frames a second.*

For further information contact Patrick S. McCormick at [pat@lanl.gov](mailto:pat@lanl.gov).

- [1] IBM Cell Broadband Engine Technology, <http://www-03.ibm.com/technology/cell/>.
- [2] P.S. McCormick et al., *IEEE Visualization 2004*, October, 171-178 (2004).
- [3] P.S. McCormick et al., *Parallel Computing*, **33**(10-11), 648-662 (2007).

#### Funding Acknowledgments

- Department of Energy, Office of Science, Office of Advanced Scientific Computing Research

# Radio-frequency Interference Cancellation for Wideband Communications

Christopher M. Brislawn, CCS-3; Karl Hill, New Mexico State Univ.

*It's like noise-canceling headphones on steroids.*

**T**hat's one way to think of the signal processing system that could emerge from a collaboration between New Mexico State Univ. (NMSU) and LANL researchers. Their goal is to cancel strong but unwanted radio-frequency signals so that weaker signals of interest can be detected.

Noise-canceling headphones do a fine job of canceling the drone of an airplane's engines so the user can enjoy music, but this system of the future will have a more complicated task. The interference might be a Navy vessel's radar, for instance, and the signal of interest might be covert communications between terrorists.

"It's the same principle but the technology has to be different," said Chris Brislawn, a LANL scientist and affiliated faculty member in NMSU's Department of Mathematical Sciences. "We're working at radio frequencies instead of audio frequencies and with signals propagating at the speed of light rather than the speed of sound."

Working with Brislawn are Mark Dunham of LANL's International, Space, and Response Division Office, Charles Creusere and Jaime Ramirez-Angulo of NMSU's Klipsch School of Electrical and Computer Engineering (ECE), Joseph Lakey of Mathematical Sciences, and two graduate students – Scott Izu (Mathematics) and Jose Luis Ruiz Chavira (ECE).

Noise-canceling headphones use a technology called active noise control. A microphone on the headphones picks up ambient noise, which is synthesized by electronic circuitry and fed into the headphones 180 degrees out of phase to cancel the noise. "At audio frequencies this is a tried and true method," Lakey said. "We're doing this at much higher frequencies and we're trying to do it in real time, which means we have to process and feed back our signal fast enough to make a difference."

Another complicating factor, Brislawn said, is that "we can't make any assumptions about what the signal of interest might be. We don't want to prejudice our methods in favor of any particular type of signal." For instance, Fig. 1 shows a loud FM signal (synthesized in the computer) towering over a recorded urban RF background. The interference-cancellation algorithm is able to suppress this unwanted signal by over 40 dB without adversely affecting nearby potential signals of interest.

One subtlety the research must address is the fact that the amplitude of an interference signal may vary with time. Figure 2 shows a simulated FM signal whose amplitude is fluctuating on a time scale measured in microseconds. In real life this could be due to multipath interference created when a radio wave reflects off of buildings, roadways, or nearby hills, and the cancellation algorithm must accurately track the resulting variations in interference amplitude.

The researchers are working with frequencies up to 2 or 3 GHz. Cell phones, satellite broadcasts, wireless networks, garage door openers, and baby monitors all operate in that range. "One problem is that there are so many different types of signals," Izu said, "so we need to model a lot of different types of signals in order to cancel them."


Most of the work so far has been developing algorithms and computer simulations, Creusere said. "Until we know if this works in simulation, it doesn't make sense to build hardware, which would probably require designing custom integrated chips. Silicon's not fast enough, so we'll be looking at using exotic semiconductors to handle the speeds."

Ultimately such a system could have commercial as well as defense applications. "The spectrum is getting pretty crowded, with more and more devices causing interference," Lakey said. "Probably the biggest potential application would be to allow multiple devices to operate in the same part of the spectrum."


**For further information contact Chris Brislawn at [brislawn@lanl.gov](mailto:brislawn@lanl.gov).**

## Funding Acknowledgments

- Department of Defense
- Joint memorandum of understanding (MOU) with New Mexico State University and Los Alamos National Laboratory



*Fig. 1. Radio-frequency power spectrum showing a high-power FM interference signal before targeted cancellation (red) and after (blue).*



*Fig. 2. Output magnitude of the interference-cancellation algorithm (red) accurately tracking an interference signal with a time-varying amplitude envelope (blue).*

## A Recent Discovery in Anomaly Detection

Don Hush, Patrick Kelly, Ingo Steinwart, Clint Scovel, CCS-3

Anomalies are often described as rare or unusual events. This notion can be represented mathematically by defining anomalies to be points with low probability density value. Detecting anomalies in practice is hard because we don't know the probability density function, and therefore must design a detector based on incomplete information. In particular, we are given empirical information consisting of a finite collection of points sampled from the probability density function and then asked to design a detector that makes predictions (guesses) about which points are anomalous. The fraction of incorrect predictions (on both current and future points) is called the error rate, and our goal is to design a detector whose error rate is as small as possible.

One of the most fundamental and longstanding issues in anomaly detection is that there is no way to obtain a reliable estimate of the error rate for a detector, regardless of how it is designed. This means that there has been no way to quantitatively compare two (or more) detectors to determine which one is best. Researchers in CCS-3 have made a recent discovery that resolves this situation. In particular they have discovered a performance measure that is calibrated with respect to the error rate, which means that it can be used as a replacement for the error rate. Furthermore, this new performance measure can be reliably estimated in practice. Consequently it is now possible to make quantitative evaluations of performance, and to objectively answer the question, "Which of these detectors is more accurately identifying the anomalous points?" Furthermore, this discovery has opened the door to a whole new class of design methods that work by optimizing this new performance measure. In particular, CCS-3 researchers have invented a new design method that is guaranteed to be computationally efficient and to provide near-optimal results for nearly all probability densities encountered in practice.

For further information contact Don Hush at [dhush@lanl.gov](mailto:dhush@lanl.gov).

### References:

D. Hush and J. Howse, "Anomaly Detection on Graphs," Los Alamos National Laboratory Technical Report LA-UR-05-8440 (2005).

D. Hush, P. Kelly, C. Scovel and I. Steinwart, "Provably Fast Algorithms for Anomaly Detection," Los Alamos National Laboratory Technical Report LA-UR-05-4367 (2005).

I. Steinwart, D. Hush, and C. Scovel, *Journal of Machine Learning Research*, **6**, 211-232 (2005).


C. Scovel, D. Hush, C. Scovel and I. Steinwart, *Analysis and Applications*, **3**(4), 356-371 (2005).

I. Steinwart, D. Hush, and C. Scovel, *Neural Information Processing Systems*, **17**, 1337-1344 (2005).

### Funding Acknowledgments

- Los Alamos National Laboratory Directed Research and Development Program





*Fig. 1. An anomaly detected on Mount Rushmore.*

# Data Mining in Radiation Portal Monitoring

Tom Burr, Jim Gattiker, Kary Myers, CCS-6; George Tompkins, D-6

Currently deployed passive gamma and neutron detectors screen for illicit nuclear material. Archived data can help to evaluate detection probabilities (DP) and to investigate several issues, including: 1) background gamma suppression, 2) nuisance gamma alarms arising from naturally occurring radiation (NORM), and 3) the state of sensor health.

Figure 1 shows one screening location, where four detector panels each record a neutron and a low- and high-energy gamma count every 0.1 s for 5 to 20 s, resulting in a 12-component time series of 50 to 200 observations.

Figure 2 (top left) is an example of the raw and smoothed scaled low-energy gamma count time series (“profile”) from one vehicle. Figure 2 (top left) illustrates background suppression, which modeling suggests arises from displacement of the air from which background gammas arise, and from shielding nearby ground sources such as asphalt. Threat detection algorithms that rely on anticipated profile shapes must consider the effect of background suppression. Several options to adjust for background suppression have been evaluated [1-3]. Note from Fig. 2 (bottom left) that simply subtracting the average background suppression (the “template,” with alignment to adjust for unequal profile lengths) results in undesirable patterns in the residuals. One advantage of monitoring count ratios is that their suppression is less (right side plots in Fig. 2).

Nuisance alarms due to NORM limit DP for threats. Strategies to recognize common NORM such as cat litter or ceramics depend on the sensor energy resolution. One of the best methods using the systems described here (two-energy gamma and neutron) uses a nonparametric density estimation method for pattern recognition [4-5].

Sensor health can be monitored using periodic check-source measurements, but because the unshielded background changes over time due to environmental changes, archived data is a potential quality control (QC) component to flag measurement anomalies. One QC option is to monitor count ratios. For example, using training data from December 1-15, 2003, from one site, and testing data from January-March of 2005, a nominal 1% false alarm rate derived from selecting ratio alarm thresholds from the training data had an actual false alarm rate of 1 to 40% in the testing data, thus indicating nonstationarity [6].

For further information contact Tom Burr at [tburr@lanl.gov](mailto:tburr@lanl.gov).

- [1] J. Gattiker and T. Burr, to appear in *Journal of Nucl. Mater. Management* (2008).
- [2] T. Burr et al., *Applied Radiation and Isotopes*, **65**, 569–580 (2007).
- [3] T. Burr et al., “Statistical Evaluation of the Impact of Background Suppression on the Sensitivity of Passive Radiation Detectors,” in *Statistical Methods in Counter Terrorism*, New York: Springer (2006).
- [4] T. Burr and K. Myers, to appear in *Applied Radiation and Isotopes* (2008).
- [5] T. Burr and J. Doak, *Intelligent Data Analysis*, **11**, 651–662 (2007).
- [6] P. Doctor et al., “Evaluation of State of Health Monitoring for Radiation Portal Monitors,” PNNL-SA-4875 (2006).

## Funding Acknowledgments

- Department of Homeland Security



Fig. 1. Example screening location with four detector panels surrounding the vehicle.

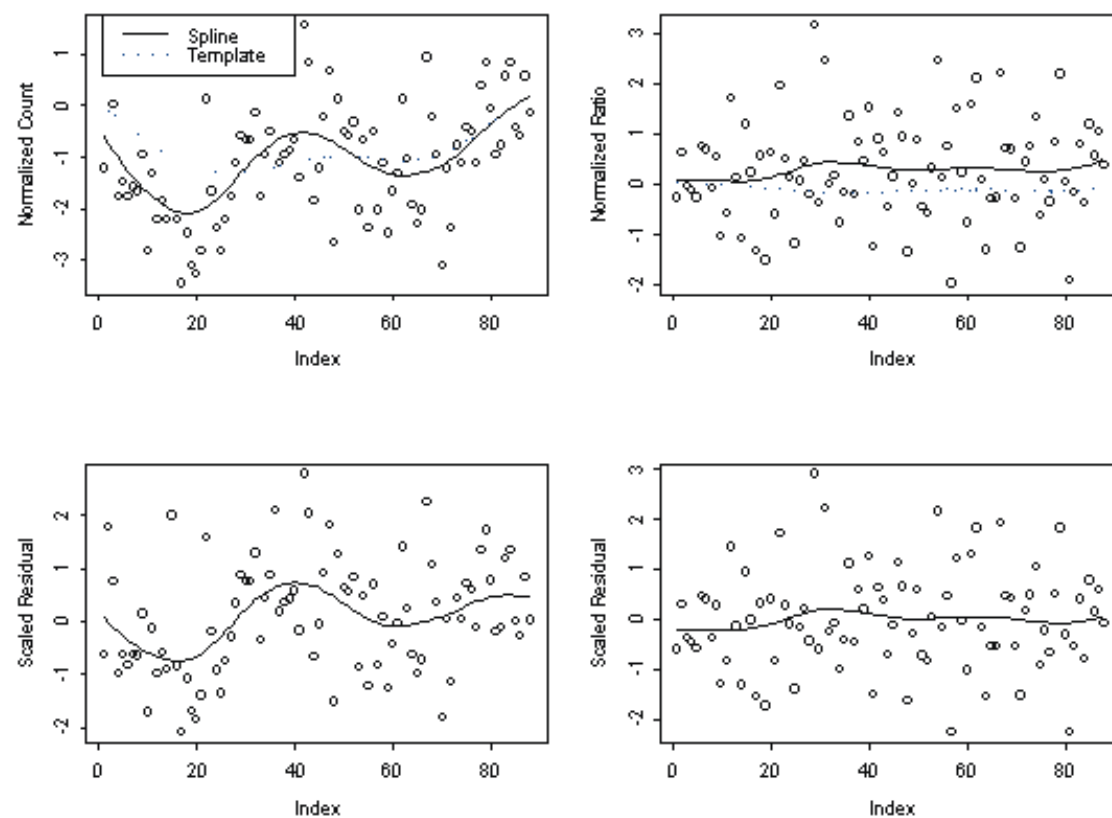


Fig 2. (top left) Example scaled low-energy gamma profile from one vehicle, a spline fit, and an average ("template") over many vehicles; (top right) same as top left, but for the gamma count ratio (defined as low-energy gamma count)/total gamma count); (bottom left) scaled residual from the count, and (bottom right) scaled residual from the count ratio.

## Framework for Responding to Network Security Events (FRNSE)

Justin “J.D.” Doak, Aaron Lovato, CTN-5

**L**ANL network security analysts face a difficult challenge in monitoring the high volume of logs from network sensors around the Laboratory and taking action on potential security threats. The purpose of the FRNSE (pronounced fren-zee) project is to automate the process of monitoring log files and sensor output and to then respond in near real-time to potential security threats. FRNSE integrates sensors, policy, and responses in a centralized manner providing consistency in our approach to network security operations.

The FRNSE architecture (Fig. 1) consists of three main components: sensors, agents, and one or more servers. Each component plays a critical role in detecting and responding to network security threats. Sensors are typically network appliances (e.g., TippingPoint) that analyze packets crossing the section of the network where they are connected. Some of these appliances are in-line and have the ability to block network traffic. These tools are referred to as intrusion prevention systems (IPS). Other appliances (e.g., Snort) are not in-line and can only detect potentially malicious behavior. These tools are referred to as intrusion detection systems (IDS). In addition to looking at raw packets, certain sensors look at flow data, which summarizes packet transfers between two hosts with an established connection. Sensors are the first link in the FRNSE response pathway because they alert the system that something of interest may have occurred.

Agents act as an intermediary between the sensors and the server(s). Sensors may generate a large number of alerts; however, we may only be interested in a small subset of those alerts. An agent is configured to monitor the output of a sensor and only generates FRNSE alerts (alerts sent to the FRNSE server) when the sensor produces output that we have predetermined to be of interest. Agents formulate a valid FRNSE alert from the raw sensor output and send that message to the server(s) over a secure socket layer (SSL) connection. Alerts contain a category, a severity, and a certainty, among other fields. Category is the type of alert generated (e.g., ExploitAttempt), severity is how damaging an event is, and certainty is our confidence that the event for which we are generating an alert actually occurred

(i.e., it is not a false positive). These three fields, when combined with a policy, determine how we respond to an alert, if at all. See Fig. 2 for an example of a FRNSE alert.

The FRNSE server(s) provides a centralized mechanism for integrating the various network security appliances and responding to their output in a consistent manner. When alerts are first received by the FRNSE server, they undergo correlation processing. Correlation uses previous alerts and any other pertinent data to modify fields in the alert, such as increasing or decreasing severity, or perform other actions. For example, if we have seen different categories of alerts generated by the same IP address in the last 24 hours, we increase the severity. After correlation, we respond according to policy. Available responses include switch blocks, firewall blocks, ticket posting, vulnerability registration, and e-mail notification. The results of those responses are obtained, and then the alert, along with the responses and

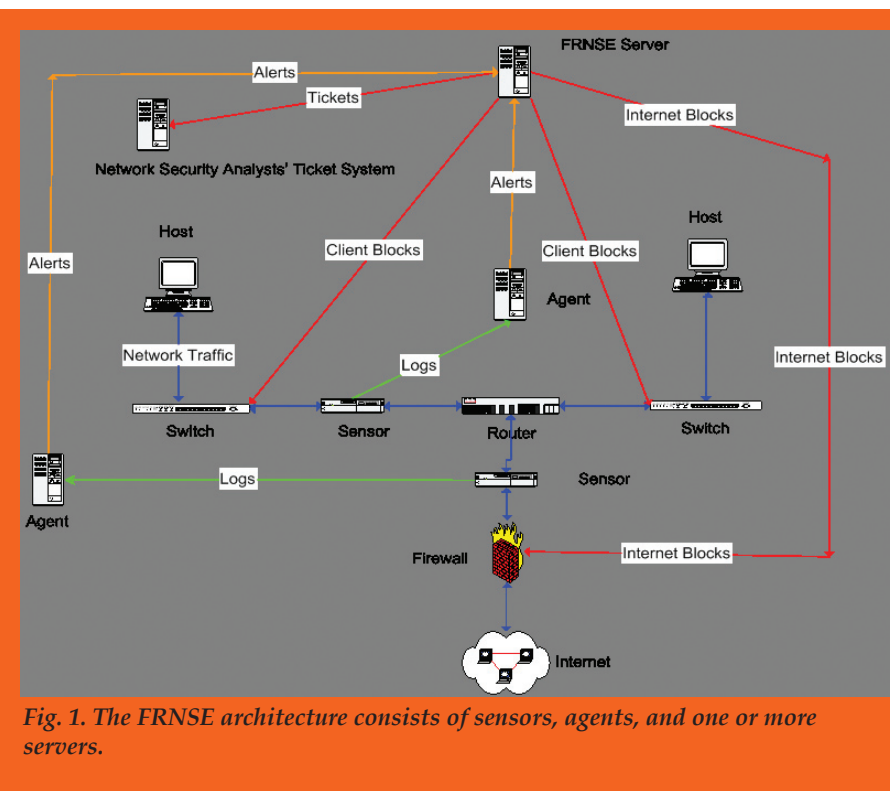


Fig. 1. The FRNSE architecture consists of sensors, agents, and one or more servers.



```

<alert>
  <category>ExploitAttempt</category>
  <severity>1.0</severity>
  <sensor>CFlow</sensor>
  <certainty>1.0</certainty>
  <taskmsg>srcip = 128.165.251.231 count = 11 start.date = 2007-11-28 start.ctime =
    15:03:06 end.ctime = 15:03:06 clock = 1993812</taskmsg>
  <srcip>128.165.251.231</srcip>
  <usermsg>Worm activity (RPC related) was detected on this host. Please visit
    http://patchme.lanl.gov/</usermsg>
  <type>CFlow</type>
</alert>

```

Fig. 2. The XML-style format of a FRNSE alert.

results, is stored in a database. The FRNSE GUI (Fig. 3) is an interface built on the alert database to provide analysts with access to the raw alert data.

To illustrate the operation of FRNSE, let's follow the life of a FRNSE alert from creation at the sensor to a set of responses in the FRNSE server. Let's say that a rule for one of the Snort sensors is triggered by a particular network packet, generating an entry in a log file. The Snort Agent, which is monitoring that log file, determines that this is one of the Snort rules for which we generate FRNSE alerts. This alert is sent over SSL to the FRNSE server where it undergoes correlation. The severity is doubled in this correlation step due to historical alerts generated by this IP. The category, severity, and certainty of the alert are then compared with policy to determine how to respond. We respond by requesting a switch block and posting a ticket for the network security analysts.

FRNSE has been actively responding to network security events since November 2006. It has processed millions of alerts that then generated tens of thousands of responses. (Many alerts are archived, but not responded to.) All of this is performed in near real-time so that potential security

threats are rapidly mitigated before spreading and causing more damage. Possible future work includes tuning sensors to detect more threats, developing agents for more sensors, adding new responses (e.g., e-mail blocking), and improving correlation.

For further information contact Justin "J.D." Doak at [jdoak@lanl.gov](mailto:jdoak@lanl.gov).

### Funding Acknowledgements

- Computing, Telecommunications, and Networking Division (CTN) Classified and Unclassified Computer Security Programs

**FRNSE v0.3 Client**

Real-Time Status Settings

**Real-Time Alerts**

Clear Options Status: Records current as of 2007-07-26 14:30:57.717.

Actions	Ticket	Alert ID	Start TS (MT)	Sensor	Category	Severity	Certainty	Message
Remove		783537	2007-07-25 11:12:21.0	TippingPoint	BlockedScan	Med (0.500)	N/A	Invalid TCP Traffic: Possible nmap Scan (SYN FIN)
Remove		771287	2007-07-23 08:52:00.0	EmaadFlows	Scan	Med (0.451)	N/A	This host had anomalous activity on port udp/123. Contact
Remove		766399	2007-07-22 06:24:38.0	TippingPoint	BlockedExploitAttempt	High (0.750)	N/A	blocked for Sasser Virus Infection
Remove		761142	2007-07-20 16:59:47.0	TippingPoint	BlockedScan	Med (0.500)	N/A	Invalid TCP Traffic: Possible nmap Scan (No Flags)
Remove		761138	2007-07-20 16:57:58.0	TippingPoint	BlockedScan	Med (0.500)	N/A	Invalid TCP Traffic: Possible nmap Scan (No Flags)
Remove		760681	2007-07-20 14:47:38.0	TippingPoint	BlockedScan	Med (0.500)	N/A	Invalid TCP Traffic: Possible nmap Scan (SYN FIN)
Remove		760618	2007-07-20 14:33:15.0	TippingPoint	BlockedScan	Med (0.500)	N/A	Invalid TCP Traffic: Possible nmap Scan (SYN FIN)
Remove		756543	2007-07-19 15:30:00.0	Sygate	BlockedExploitAttempt	Med (0.600)	N/A	[474] HTTP ACF File Parsing Buffer Overflow Attempted (9
Remove	627	749110	2007-07-18 07:14:00.0	EmaadFlows	Scan	Med (0.529)	N/A	This host had anomalous activity on port udp/14609. Conta
Remove		742759	2007-07-17 02:50:55.0	TippingPoint	BlockedScan	Med (0.500)	N/A	Invalid TCP Traffic: Possible nmap Scan (SYN FIN)
Remove	619	734745				.925)	N/A	This host had anomalous activity on port udp/1501. Conta
Remove	605	718167		EmaadFlows, Scan		.627)	N/A	This host had anomalous activity on port udp/123. Contact
Remove	604	714295		Start Time	2007-07-18 07:14:00.0	.600)	N/A	Somebody is scanning your computer. Your computer's UI
Remove	603	714259		Duration	0	.600)	N/A	Somebody is scanning your computer. Your computer's UI
Remove	600	708212		Severity	Med (0.529)	.594)	N/A	This host had anomalous activity on port udp/21523. Conta
Remove	599	701695		Certainty	N/A	.652)	N/A	This host had anomalous activity on port udp/6502. Conta
Remove		691920		Source IP	128.165.38.31	.750)	N/A	This system attempted to exploit the MS 06-040 Server Se
Remove		691919		DISARM	Get Records	.750)	N/A	This system attempted to exploit the MS 06-040 Server Se
Remove		675900		LFAP	Get Flows	.750)	N/A	This system attempted to exploit the MS 06-040 Server Se
Remove		675901		Message	This host had anomalous activity on port udp/14609. Contact CSIRT at 5-8641	.750)	N/A	This system attempted to exploit the MS 06-040 Server Se
Remove		675899		Result	post - 2007/07/18 07:14:00 Scan, ticket #627	.750)	N/A	This system attempted to exploit the MS 06-040 Server Se
Remove	590	670363		Ticket	627	.000)	N/A	This host had anomalous activity on port icmp/0. Contact C
Remove	586	669952				.000)	N/A	This host had anomalous activity on port tcp/80. Contact C
Remove		669702				.000)	N/A	This host had anomalous activity on port icmp/0. Contact C
Remove	585	669582	2007-07-03 10:26:00.0	EmaadFlows	Scan	Critical (1.000)	N/A	This host had anomalous activity on port tcp/80. Contact C

Fig. 3. The web-based FRNSE GUI provides an interface to the FRNSE alert database.



# Combining MEG and fMRI for Single-pass MEG Analysis – A Simulation Study

Doug Ranken, HPC-1

In magnetoencephalography (MEG), a helmet containing a large array of highly sensitive superconducting quantum interference device (SQUID) sensors is used to measure the weak magnetic fields arising from human brain activity. In a typical MEG sensory response experiment, a single MEG pass measures the brain activity corresponding to a single application of a stimulus, such as the sounding of a tone. Usually, 100 or more passes are recorded for a given stimulus, and their average is used to perform location and timecourse analysis of the stimulus-induced brain activity. Inferring locations and timecourses of brain activity from MEG measurements, the MEG inverse problem, is a hard problem to solve, and has been addressed with a wide array of approaches. A subset of these are dipole model approaches that assume the sources of activity are point-currents (dipoles). One approach developed at LANL is the Calibrated-Start Spatio-Temporal (CSST) Dipole Model. In CSST, the number of regions of brain activity is first estimated, then nonlinear simplex searches are applied to multiple sets of starting locations to find the set of dipoles and their associated timecourses that best account for the measured data. Because of the low signal-to-noise ratio, inverse analysis of single-pass MEG data is hard, or impossible. In some cases, using information from other sources could aid in this analysis.

In functional magnetic resonance imaging (fMRI), a magnetic resonance imaging (MRI) scanner is used to detect changes in blood properties, such as oxygenation, to find locations of increased brain activity. One standard approach uses a statistical analysis of the blood oxygenation level dependent (BOLD) response, generating a T-statistic for each voxel in the fMRI volume that provides a measure of brain activity. The fMRI data in a combined MEG-fMRI experiment can complement the MEG data, in terms of providing additional brain activity location information, but does not provide useful information on activity timecourses.

The MEG Forward Simulator in MRVIEW software was used to create three regions of brain activity, with activity timecourses similar to those shown in the upper right

plot in Fig. 1, and MEG sensor timecourses shown in the two left plots. The upper left plot shows noise-free sensor timecourses; the lower left plot shows the same timecourses with added simulated noise, giving a noise level typical of single-pass MEG data.

A simulated fMRI volume was generated by applying a smoothing algorithm to the voxels of the three simulated regions, to create brain activity maps similar to those obtained in an fMRI analysis. A probability volume was obtained for these regions by scaling the voxel values in the smoothed volume to the interval  $[0,1]$ . This volume was registered to the MEG coordinate space. During the CSST search process, the voxel values of the probability volume were used to weight the MEG-only Chi-square measure for a given set of locations, to produce a combined (simulated) MEG-fMRI measure.

The lower-right plots in Fig. 1 show the benefits that may be realized with a combined MEG-fMRI analysis. The combined MEG-fMRI analysis demonstrates the fMRI influence on dipole locations. In the combined analysis, the fitted dipole locations are pulled toward regions of high fMRI activity, counteracting the influence of the MEG noise on the dipole fits. The dipole timecourses arising from the combined CSST analysis of the noisy simulated MEG data more closely match the low-noise analysis results, as evidenced by the shape and peak amplitudes shown in green and blue. In the MEG-only case, the peak green dipole timecourse amplitude is 50% higher than it should be, because of complications in dipole fitting arising from the sensor noise. A similar problem can be seen with the blue dipole. For both the green and blue dipoles, the combined analysis timecourse results closely match the actual timecourses. This improvement in timecourse matching most likely arises from the fMRI contribution reducing the error in fitted dipole locations due to the influence of sensor noise.

This probability volume approach to combined MEG-fMRI analysis has also been adapted for use in a Bayesian MEG-fMRI analysis. Probability volumes can be constructed from other types of data, such as structural MRI, positron emission tomography, or MEG analysis results obtained from averaged (low-noise) MEG data.

**For further information contact Doug Ranken at [ranken@lanl.gov](mailto:ranken@lanl.gov).**

## Funding Acknowledgments

- Los Alamos National Laboratory Directed Research and Development Program
- MIND Research Network
- National Institutes of Health

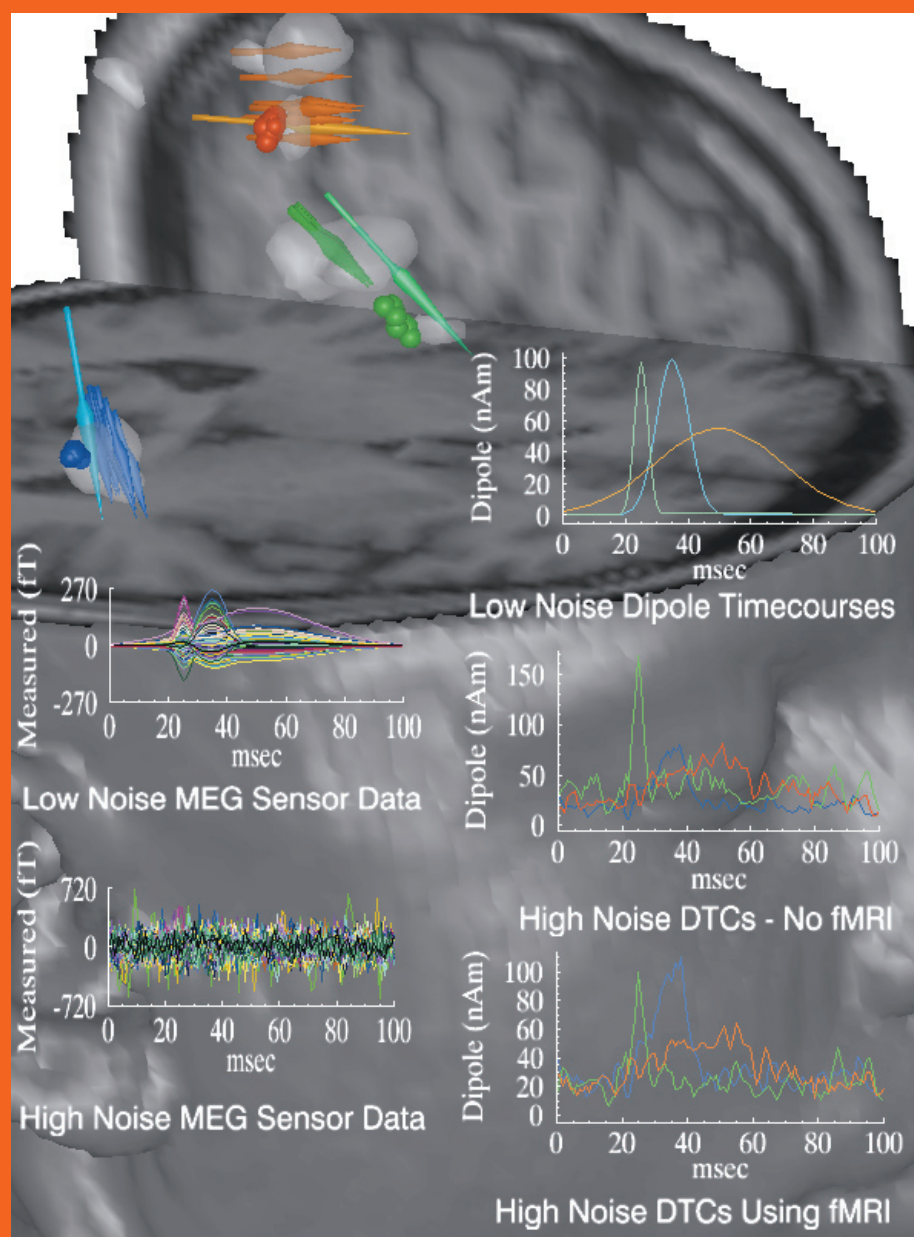


Fig. 1. The MRI head data shown in 3D was used to create simulated regions of brain activity. The white regions represent simulated fMRI data. The 3-dipole CSST solution for the low-noise case is shown with the larger arrows. For the high-noise analyses, the MEG-only solution locations are shown using spheres. The combined MEG, fMRI solution locations are shown using the smaller arrows.

# Adaptive and Dependable Server I/O Networking Support in a High Performance Cluster Computing Environment

Hsing-Bung Chen, HPC-5

We present an adaptive and dependable server I/O fault-management mechanism used in equal cost multi-path routing to enable LANL's parallel scalable back bone (PaScalBB) (Fig. 1) high performance computing (HPC) systems to run computational jobs around the clock without service interruption during unexpected physical I/O link failures and connection loss. This mechanism, named Dead Server I/O Gateway Detection and Recovery (DGD) (Fig. 2), can detect a data path connectivity problem within seconds when it happens. Then DGD removes the entry of a dead I/O gateway from a multi-path routing table, migrates a connecting I/O path to an available entrance in a multi-path routing table, and preserves and resumes the existing I/O data stream. DGD can tolerate multiple single points of failures, keep the streaming I/O data moving, and seamlessly continue and finish computation jobs. Figure 3 illustrates the self-explained pseudo-code for the proposed DGD mechanism. We have developed a proof-of-concept implementation of this proposed DGD mechanism on a Linux cluster as a blueprint for a production-type reliability-availability-serviceability (RAS) solution. Figures 4 and 5 show the testing cases of using this DGD process on BlueSteel (256-node InfiniBand cluster). Eventually we plan to apply this process to all LANL's PaScalBB-based HPC production clusters.

For further information contact Hsing-Bung Chen at [hbchen@lanl.gov](mailto:hbchen@lanl.gov).

## Funding Acknowledgments

- Department of Energy, National Nuclear Security Administration

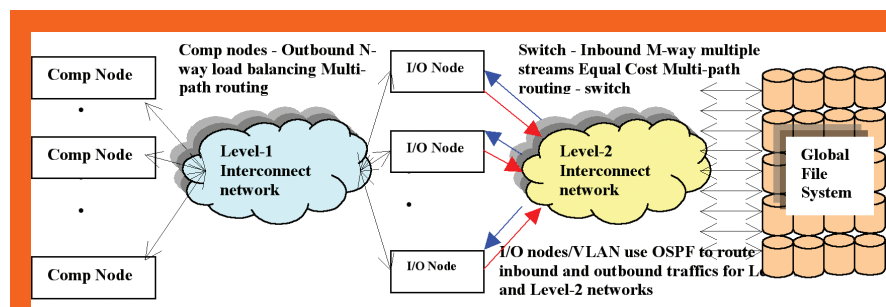


Fig. 1. Components used in LANL's PaScalBB backbone infrastructure.

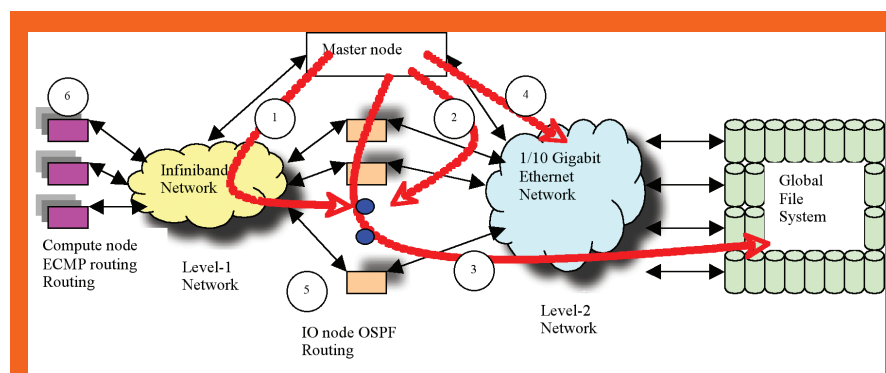


Fig. 2. DGD and ECMP route recovery system diagram and processing steps.

```

initDGD();
while (1) {
  1. Check Level-1 connectivity:
    1.1 Master node → Level 1 Network → I/O nodes (Level 1 interface)
    1.2 Master node → Level 1 Network → I/O nodes (Level 2 interface)
  2. Check Level-2 connectivity:
    2.1 Master node → Level 2 Network → I/O nodes (Level 2 interface)
    2.2 Master node → Level 2 Network → I/O nodes (Level 1 interface)
  3. Check Global File System connectivity:
    Master node → Level 1 Network interface → Global File System
  4. Check Status of the Level-2 Interconnect: Master node → Level 2 Storage network
  5. Check OSPF status on both IO-gateway and Level-2 interconnect:
    OSPF routing status : Active vs. InActive
  Summarize of step1-step5 checking and set current I/O node status
  Update I/O nodes state machine and detect ECMP route fault-events
  if (PreviousStatus(I/O node) equal CurretnStatus(I/O node)) {
    Active → Active: No ECMP route change
    InActive → InActive: No ECMP route change
    ECMP route changed → NO
  }
  if (PreviousStatus(I/O node) not equal CurretnStatus(I/O node)) {
    Active → InActive: Remove dead/inactive I/O node from ECMP route
    InActive → Active: Add new Active I/O node back to ECMP route
    ECMP route changed → Yes
  }
  if (ECMP route changed) {
    6. Adaptive Event-handling
    + Generating syslog information
    + Email DGD status updated notification
    + Record DGD history status
    + Adaptive route change on Equal Cost Multipath (ECMP) Route
    → Add or Remove I/O node from ECMP route
  }
  sleep N seconds; // kick-off DGD process every N seconds
}

```

Fig. 3. DGD processing pseudo code.




Fig. 4. Normal Write then Read with no dead I/O gateway.




Fig. 5. Write then Read with multiple dead I/O gateways involved during testing.



## Petascale Data Storage Institute

Gary Grider, HPC-DO; Carolyn Davenport, James Nunez, HPC-5

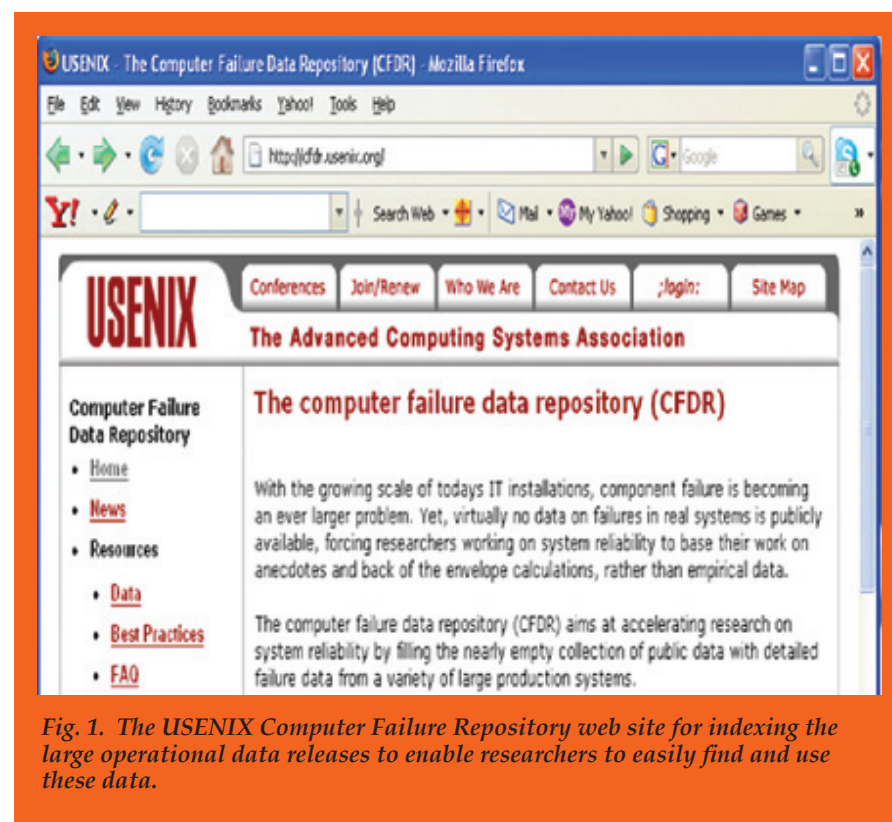
**P**etascale computing infrastructures for scientific discovery make petascale demands on information storage capacity, performance, concurrency, reliability, availability, and manageability. The last decade has shown that parallel file systems can barely keep pace with high performance computing along these dimensions, posing a critical challenge when petascale requirements are considered. The LANL Petascale Data Storage Institute (PDSI) focuses on the data storage problems found in petascale scientific computing environments with special attention to community issues such as interoperability, reliability, scaling, standards, and shared tools. Leveraging the experience in applications and diverse file and storage systems expertise of its members, the institute allows a group of researchers to collaborate extensively on developing requirements, standards, algorithms, reliability, understanding, and performance tool development. Mechanisms for petascale storage and results are made available to the petascale computing community. Additionally, the institute holds periodic workshops and develops educational materials on petascale data storage for science.

**Collaborators.** The PDSI is a collaboration between researchers at Carnegie Mellon University, National Energy Research Scientific Computing Center, Pacific Northwest National Laboratory, Oak Ridge National Laboratory, Sandia National Laboratories, Los Alamos National Laboratory, the University of Michigan, and the University of California, Santa Cruz. The Los Alamos participation in the PDSI is through the LANL Institutes office via the Information Science and Technology Institute and the High Performance Computing (HPC) Division Systems Integration Group.

**Recent Accomplishments.** LANL released 9 years of computer failure data; over 23,000 records for several thousand machines. Additionally, LANL released several million usage records (job size, processors/machines used, duration, time, etc.) for the same machines for which the failure data was released. LANL also released disk

failure data for a subset of these same machines. This data is publicly available and has received over 900 downloads in the last 6 months. The data has generated top paper awards at two conferences. Papers by Dr. Bianca Schroeder and Dr. Garth Gibson of Carnegie Mellon University won prizes at the 5th USENIX Conference on File and Storage Technologies (FAST 2007) and the International Symposium on Dependable Systems and Networks (DSN 2006).

The LANL data release and subsequent groundbreaking publications have encouraged an outpouring of operational data from a number of supercomputer sites and internet data and service providers. The USENIX, the Advanced



*Fig. 1. The USENIX Computer Failure Repository web site for indexing the large operational data releases to enable researchers to easily find and use these data.*














Computing Systems Association, has begun a project to index all of these large-scale operational data releases on the internet (Fig. 1). LANL is helping to organize and lead this effort.

Additionally, LANL released an important metadata synthetic benchmark, the parallel file tree walk application, which has received over 100 downloads. This is the first parallel file tree walker benchmark application of its kind. Additionally, LANL has cleared a process for releasing parallel I/O traces and file system statistics from hundreds of workstation archives and is poised to release several gigabytes of parallel I/O traces for use in the research community.

**For more information contact Gary Grider at [ggrider@lanl.gov](mailto:ggrider@lanl.gov).**

### **Funding Acknowledgments**

- Department of Energy, Office of Science, Scientific Discovery through Advanced Computing Program



# Materials

To some, materials science is simply the study of the properties of materials, and phase transitions therein, at macro and microscopic levels. However, at the fundamental level, materials lie at the heart of everything we do for the Nation at LANL—from creating new and efficient forms of energy production to stockpile stewardship—and live at the heart of everything we touch in our daily lives—from telecommunications to transportation. The Laboratory has an obvious deep and rich scientific history in materials science, reaching back to the origins of the Laboratory itself. It is therefore not surprising that LANL is applying its scientific acumen to the furthest reaches of materials science at the fundamental theoretical and applied levels.

Material science exists at the confluence of many disciplines, and the papers in this section reflect this diversity in both scientific discipline and application. Some topics represented in this section are:

- Breakthroughs in theory and modeling for materials at multiple scales (nano to macro), to ideas for new and efficient energy sources
- Applications of statistical methods for extracting new insights on the reliability of materials for military applications
- Development of accurate equations of state to predict material behavior in multiple environments.
- Modeling, simulation, design, manufacturing, and evolution of nuclear fuels—an important emerging strategy in the Nation's quest for energy security.
- Theoretical and modeling developments for materials and liquids germane to energy storage, retrieval, and production.

Materials science is a key historic strength of LANL, is central to the development of new technologies, and is critical to the continuing and enduring strong positive impact of LANL on society.



# Statistical Models for Stockpile and Complex System Health Assessment

Alyson Wilson, Mike Hamada, Christine Anderson-Cook, Todd Graves, Scott Vander Weil, Aparna Huzurbazar, Benjamin Sims, Greg Wilson, CCS-6

The LANL Statistical Sciences Group (CCS-6) has an ongoing applied research effort to develop methods and software tools to help assess the health of conventional and nuclear stockpiles. Given the complex technical and data issues associated with understanding stockpiles, CCS-6 works within LANL to apply the best data analysis methods to support decision making about the enduring U.S. nuclear stockpile. Likewise, CCS-6 collaborates with the Department of Defense (DoD) and other agencies on problems that improve the understanding of weapon stockpile assessment and complex system reliability.

We have developed statistical methodology for a number of problems arising in reliability assessment. This includes appropriately handling multilevel data using reliability block diagram and fault tree representations, using age and usage covariates to improve reliability estimates, combining expert knowledge with available data, and faithfully dealing with the less-than-ideal data. We have also developed and applied methodology to associate uncertainties with reliability assessments. Finally, much of the implementation has been done using GROMIT and YADAS, software tools we have developed to address the system modeling and statistical challenges often presented by weapons data [1,2] (see also [yadas.lanl.gov](http://yadas.lanl.gov)).

Our recent successes in supporting the weapons program have included the following effort:

- Combine diverse data types (e.g., pass/fail, accelerated, degradation, lifetime and specification data, and expert judgment) at multiple levels (e.g., system, subsystem, component) to evaluate system reliability [3,4,5,6].
- Model reliability as a function of multiple lifecycle covariates [7].
- Take indirect measurements to make inference about characteristics based on direct measurements that were not taken [8].
- Estimate the proportion of a population with an attribute from data that were purposely biased to contain units with that attribute [9].



*Fig. 1. CCS-6 is collaborating with Naval Surface Warfare Centers to assess the reliability of the RAM and AMRAAM missile systems.*

- Analyze degradation data with a proposed model that implies a Weibull lifetime distribution and determine how reliability can then be assessed [10].
- Support assessments of reliability margins with statistical models and analyses [11].
- Understand how to collect degradation data over time for applications like shelf-life programs [12].
- Track and trend surveillance streams from LANL stockpile weapons to help assess stockpile health [13,14,15,16,17,18].
- Enhance and expand capabilities for YADAS [14,19].

Past and current DoD collaborators include AMRDEC Huntsville, MCPD Fallbrook, NSWC Corona, and NSWC Yorktown.

**For more information contact Christine Anderson-Cook at [c-and-cook@lanl.gov](mailto:c-and-cook@lanl.gov).**






Fig. 2. Andrew Wiedlea and Nick Hengartner examine a test set used to collect data on the health and functionality of components from weapons like the MK 67 submarine-launched mine in the picture.




Fig. 3. CCS-6 is working with the LANL Core and Enhanced Surveillance programs to track and trend weapon surveillance data streams. The goal is to assess key weapon characteristics against known specifications and estimate the distribution of these characteristics in the stockpile now and in the future.

- [1] A. Koehler, R. Klamann, "GROMIT: a tool for systems ethnography," Los Alamos National Laboratory report LA-UR-04-4370.
- [2] T. Graves, "YADAS: An Object-Oriented Framework for Data Analysis Using Markov Chain Monte Carlo," Los Alamos National Laboratory report LA-UR-01-4804.
- [3] T.L. Graves, M. Hamada, "Bayesian methods for assessing system reliability: models and computation," in *Modern Statistical and Mathematical Methods in Reliability*, World Scientific Publishing Company, A. Wilson et al., Eds., 41-54 (2005).
- [4] A. Wilson et al., *Stat. Sci.*, **21** (4), 514-531 (2006).
- [5] C. Anderson-Cook, et al., *Mil. Oper. Res.*, **12** (2), 25-37 (2007).
- [6] C.S. Reese et al., "A hierarchical model for the reliability of an anti-aircraft missile system," Los Alamos National Laboratory report LA-UR-05-9281. Under revision for *J Am. Stat. Assoc.*
- [7] C. Anderson-Cook, A. Koehler, E. Lawrence, "Model of multiple lifecycle variables," Los Alamos National Laboratory report LA-CP-07-0199.
- [8] T.L. Graves, M. Hamada, *Qual. Eng.*, **17**, 555-559 (2005).
- [9] T.L. Graves, M. Hamada, *Qual. Reliab. Eng. Int.*, **22**, 385-392 (2006).
- [10] M. Hamada, *Qual. Eng.*, **17**, 615-620 (2005).
- [11] J. Booker et al., *Mil. Oper. Res.*, **11**, 27-44 (2006).
- [12] E. Kelly, S. Vander Wiel, "Planning and evaluating sampling designs to determine aging effects," Los Alamos National Laboratory report LA-UR-06-7244.
- [13] C. Faust et al., "Analysis of Certification and Surveillance Data for Pits: Report 1," Los Alamos National Laboratory report LA-CP-07-0090.
- [14] T. Graves et al., *Technometrics*, **49**, 164-171 (2007).
- [15] M. Hamada, et al., "Analysis of Certification and Surveillance Data for Pits: Report 3," Los Alamos National Laboratory report LA-CP-07-0907.
- [16] M. Hamada et al., "Analysis of Certification and Surveillance Data for Pits: Report 2," Los Alamos National Laboratory report LA-CP-07-0908.
- [17] M. Hamada, et al., "Analysis of Certification and Surveillance Data for Pits: Report 4," Los Alamos National Laboratory report LA-CP-07-1440.
- [18] M. Hamada, et al., "Analysis of Certification and Surveillance Data for Pits: Report 5," Los Alamos National Laboratory report LA-CP-07-1441.
- [19] T.L. Graves, *J. Comput. Graph. Statist.*, **16**, 24-43 (2007).
- [20] T. Graves et al., *Technometrics*, **49**, 164-171 (2007).

#### Funding Acknowledgments

- Joint memorandum of understanding (MOU) between Department of Energy and Department of Defense



## Models and Simulations of $\text{UO}_2$ Fuel Materials Properties

Marius Stan, CCS-2; Sven Rudin, John Wills, T-1; Blas P. Uberuaga, Steven M. Valone, MST-8; Shenyang Hu, Pacific Northwest National Laboratory; Petrica Cristea, Univ. of Bucharest

Nuclear fuels are subject to severe radiation environments and their thermal, chemical, and mechanical properties change significantly with time and irradiation level [1]. The major factors that influence the properties are temperature, stoichiometry, and microstructure (especially porosity and point defects). In particular, the accumulation of fission products in gas bubbles can decrease the heat transfer, leading to overheating of the fuel element and local melting (Fig. 1).

This work is focused on the irradiation effects on properties such as thermal conductivity, oxygen diffusivity, and thermal expansion. Our methods cover a large spectrum of time and space scales, from electronic structure to atomistic, to meso-scale, to continuum (Fig. 2).

The results include electronic structure calculations of phonon spectra (Fig. 3), models of point defect concentrations, stoichiometry and diffusivity (Fig. 4), simulations of gas bubble formation and evolution (Fig. 5), and coupled simulations of heat transport, diffusion, and thermal expansion [6]. Future work will be focused on advanced, multi-component fuels that contain transuranic elements and fission products.

For further information contact Marius Stan at [mastan@lanl.gov](mailto:mastan@lanl.gov).

- [1] D.R. Olander, "Fundamental Aspects of Nuclear Reactor Fuel Elements," TID-26711-P1, Technical Information Service, U.S. Department of Commerce, Springfield, Virginia.
- [2] M.F. Lyons, et al., *Trans. Amer. Nucl. Soc.*, **8**, 42 (1965).
- [3] M. Stan, et al., *J. Alloys Comp.*, **444–445**, 415–423 (2007).
- [4] P. Cristea, M. Stan, and J. C. Ramirez, *J. Optoelectr. Adv. Mater.*, **9**, 1750-1756 (2007).
- [5] S.Y. Hu, et al., *Appl. Phys. Lett.*, **90**, 81921-81923 (2007).
- [6] J.C. Ramirez, M. Stan, and P. Cristea, *J. Nucl. Mater.*, **359**, 174-184 (2006).

### Funding Acknowledgments

- Department of Energy Global Nuclear Energy Partnership
- Department of Energy, Office of Nuclear Energy




Fig. 1. Cross-section of  $\text{UO}_2$  fuel rod operated at a linear power high enough to cause extensive melting [2].

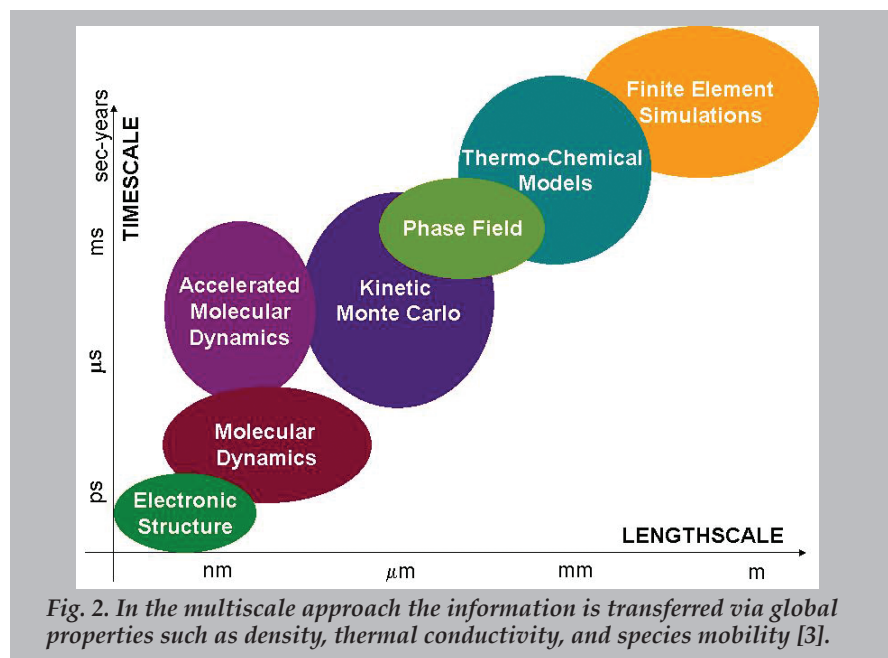


Fig. 2. In the multiscale approach the information is transferred via global properties such as density, thermal conductivity, and species mobility [3].




Fig. 3. Electronic structure calculations of phonon frequencies of  $\text{UO}_2$ . This is a complex material with properties determined by temperature, magnetic structure, strong spin-orbit coupling, and strongly correlated 5f electrons. The frequencies are used to determine thermal expansion and heat capacity of  $\text{UO}_2$ . Magnetic structure affects calculated energies and phonon frequencies significantly. These results indicate that magnetic moments must be included with specific structure [3].




Fig. 4. Oxygen self-diffusivity as a function of non-stoichiometry ( $x$  in  $\text{UO}_{2+x}$ ) and temperature. Solid line: our model [4]. Symbols: data from literature [4]. The model we have developed is valid for a large domain of temperature ( $300 \text{ K} < T < 1800 \text{ K}$ ) and oxygen content ( $x < 0.1$ ) [4].




Fig. 5. Phase Field simulation of homogeneous nucleation and growth of gas bubbles. Characteristic time: (a) 10 s, (b) 50 s, (c) 100 s, (d) 150 s, (e) 250 s, and (f) 350 s. Next step is including microstructure features, especially grain boundaries [5].

# A 3D Uranium Casting Simulation from Preheat to Cooldown

Erin I. Barker, Sharen Cummins, CCS-2

The Telluride project is developing a highly parallel, coupled-continuum scale, multi-physics simulation tool (*Truchas*). Key applications include metal-alloy casting processes of interest to the LANL foundries and development of a metal fuel-casting process for the Global Nuclear Energy Partnership (GNEP). The metal-alloy casting typically involves preheating the mold assembly by induction coils, melting the metal in the crucible, delivering the molten metal via gravity pour, and cooling to room temperature via conduction and radiation. Recently, the Telluride project has attained the capability to model this entire casting process, from initial preheat to final cooldown [1].

The numerical simulation is completed in corresponding stages: heatup, fill, and cooldown. The simulation begins by preheating the assembly, shown in Fig. 1 (a). The fill stage is modeled on a smaller mesh containing only the graphite molds and vacuum using the final enthalpies from the heatup stage as initial conditions. The uranium is introduced through an inlet velocity and temperature boundary condition maintaining the mass inflow rate from the casting [Fig. 1 (b)]. Finally, the mold is allowed to cool via radiative and conductive exchange to the walls of the vacuum chamber. As the uranium cools down, it transforms from liquid to the solid  $\Gamma$  phase at temperature  $T=1406$  K,  $\Gamma$  to  $\beta$  at  $T=1048$  K, and  $\beta$  to  $\alpha$  at  $T=941$  K.

The cooldown stage is modeled without and with thermo-mechanical deformation of the mold and cast [Fig. 1 (c)]. Enthalpies and volume fractions from the fill stage, along with the enthalpies from the heatup stage, are mapped to a full mesh. For both simulations, heat conduction, view-factor radiation, and isothermal phase change are calculated. The simulation with thermo-mechanical deformation is calculated on a smaller mesh similar to the fill stage mesh. This mesh contains gap elements to allow for gap formation and sliding between the uranium and the mold.

Simulation results were compared with thermocouples located inside and outside the mold for each stage of casting. Comparisons were used for validation of the simulation results, sensitivity studies of parameters such as conductivity and emissivity of the insulating felt (Fig. 2), and insight into including effects such as thermo-mechanical deformation during cooldown (Fig. 3).

For further information contact David A. Korzekwa at [dak@lanl.gov](mailto:dak@lanl.gov).

[1] S. Cummins and K. Lam, "An Overview of a 3D Plutonium Casting Simulation: From Mold Preheat to Cooldown," Los Alamos National Laboratory Report LA-UR-06-0751.

## Funding Acknowledgments

- Department of Energy, National Nuclear Security Administration, Advanced Simulation and Computing Program

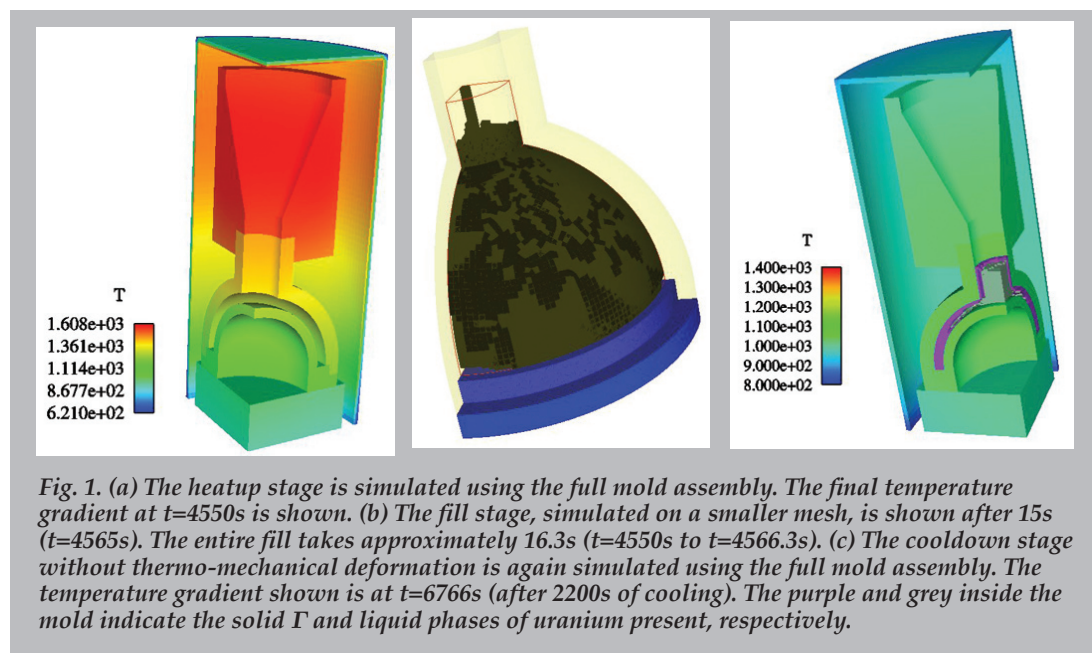


Fig. 1. (a) The heatup stage is simulated using the full mold assembly. The final temperature gradient at  $t=4550$ s is shown. (b) The fill stage, simulated on a smaller mesh, is shown after 15s ( $t=4565$ s). The entire fill takes approximately 16.3s ( $t=4550$ s to  $t=4566.3$ s). (c) The cooldown stage without thermo-mechanical deformation is again simulated using the full mold assembly. The temperature gradient shown is at  $t=6766$ s (after 2200s of cooling). The purple and grey inside the mold indicate the solid  $\Gamma$  and liquid phases of uranium present, respectively.




Fig. 2. Sensitivity of the temperature profiles to the conductivity of the insulating felt for the outer mold is shown. Conductivity values include  $K_{\text{felt}} = 0.01, 0.05, 0.1, 1.0$ . The solid line is the thermocouple data extracted from the casting. The range of results highlights the need to better understand such parameters.




Fig. 3. Temperature profile for thermocouple located on the outer model compared with the simulation results with and without thermo-mechanical deformation. The results without thermo-mechanical deformation clearly show the iso-thermal phase transitions that are inconsistent with the thermocouple data. This discrepancy highlights the importance of including such effects.



# Simulations of Heat and Oxygen Transport in a Nuclear Fuel Element

Marius Stan, CCS-2; Bogdan Mihaila, MST-6; Alek Zubelewicz, T-3; Juan C. Ramirez, EXPONENT Inc.; Petrica Cristea, Univ. of Bucharest

Uranium oxide ( $\text{UO}_{2+x}$ )-based nuclear fuels are commonly used in thermal, light-water reactors and have been recently considered as the potential fuel for fast breeder reactors. The oxide nuclear fuel rods consist of oxide fuel pellets stacked in a cylindrical metal cladding and then bundled in a fuel assembly (Fig. 1) [1] operating at temperatures of up to 2000 K. In this work we examine the influence of temperature and stoichiometry changes on the  $\text{UO}_{2+x}$  fuel properties and on the coupling of heat and species transport in a fuel element with stainless steel cladding. The objective is to improve the understanding of fuel damage and performance.

In a nuclear reactor, irradiation induces changes in material properties such as microstructure, density, thermal conductivity, specific heat, and oxygen diffusivity. We have developed models of these properties that provide correlations with temperature, pressure, burnup, and other reactor parameters [2,3]. As opposed to common approaches, our models include the dependence of the properties on the stoichiometry  $x$  in  $\text{UO}_{2+x}$ .

The finite element simulations of coupled heat and oxygen transport were performed using COMSOL Multiphysics®, which provides an ideal tool for studying coupled phenomena and allows for mesh and time step refinement in 3D configurations (Fig. 2). To solve the heat and species equation, we employed quadratic Lagrange elements and a nonlinear iterative technique with a nested unsymmetric multifrontal (UMFPACK) linear solver.

Our steady-state parametric studies were focused on determining the centerline temperature in the fuel rod as a function of non-stoichiometry and the rate of heat generation during fission. Given the strong temperature gradients in the reactor, we have included the effect of the thermally driven diffusion of species, also known as the Soret effect. We proved that the counterbalancing of the Soret and Fickian fluxes is responsible for the variation of oxygen concentration in the fuel pellet [4]. The simulations show that including the dependence of thermal conductivity

and density on non-stoichiometry can lead to changes in the calculated centerline temperature and thermal expansion displacements that exceed 5% (Fig. 3).

We have also simulated transient regimes and examined the time lag in the response of the temperature and non-stoichiometry distributions with respect to sudden changes in heat generation rate intensity and oxygen removal rate. Future work will include studies on the effects of porosity and simulations of fuel-cladding interactions.

For further information contact Marius Stan at [mastan@lanl.gov](mailto:mastan@lanl.gov).

- [1] United States Nuclear Regulatory Commission, Emergency Preparedness web page: <http://www.nrc.gov/about-nrc/emerg-preparedness/images/fuel-pellet-assembly.jpg>
- [2] M. Stan, P. Cristea, *J. Nucl. Mater.*, **344**, 213-218 (2005).
- [3] M. Stan, et al., *J. Alloys Comp.*, **444-445**, 415-423 (2007).
- [4] J. C. Ramirez, M. Stan, and P. Cristea, *J. Nucl. Mater.*, **359**, 174-184 (2006).

## Funding Acknowledgments

- Department of Energy, Global Nuclear Energy Partnership
- Department of Energy, Office of Nuclear Energy




Fig. 1. Uranium oxide fuel rod and assembly [1].






Fig. 2. Representative "slice" of the computational domain, showing a 22.5 deg. angular sector of the fuel element. The mesh was refined at the gap, where temperature and stoichiometry gradients are steeper.




Fig. 3. Top: Steady-state temperature (blue) and non-stoichiometry (green) distribution in radial direction for a  $\text{UO}_{2+x}$  fuel rod. Note the significant changes introduced by including the stoichiometry dependence of the thermal conductivity in the simulation (solid lines) compared to temperature dependence only (dashed lines). Bottom: Thermally induced expansion (displacement) of the fuel element. Note the changes introduced by including the stoichiometry dependence and the position of the maximum total displacement inside the fuel pellet. This is consistent with experimentally observed angular cracks.

# The Analysis of Self-diffusion and Migration of Spheres in Nonlinear Shear Flow Using a Traction-corrected Boundary Element Method

Alan Graham, INST-OFF; Shihai Feng, MST-7; Marc Ingber, Univ. of New Mexico;  
Lisa Mondy, Sandia National Laboratories

The goal of this research is to develop an innovative numerical simulation capability for the flow of suspensions of particles in liquids, which is relevant to several important energy-related technologies. We will incorporate effects that span time and length scales associated with molecular to macroscopic phenomena. As part of this goal, at the macroscopic scale, we need to determine how to efficiently model suspensions in situations in which the molecular-scale effects do not dominate, and to determine the limits of applicability of these models. We have developed a traction-corrected boundary element method (TC-BEM) that couples macroscale hydrodynamic forces with fine-scale physics between particles. We have shown that the TC-BEM is extremely accurate in predicting particle trajectories and can be used to calculate both particle self-diffusivity and newly defined migrations diffusivity for dilute suspensions. This work focuses on development of robust numerical simulation capabilities for suspensions of small solid particles in liquids, incorporating effects spanning diverse time and length scales. In the current research, a new method is developed to incorporate the near-field effects into the boundary element method. Rather than working with lubrication forces as was done previously [1-4], asymptotic solutions have been derived for traction fields in the interstitial regions between particles. The calculation of the traction unknowns, for a single boundary element centered around the point of nearest contact, is based on the relative motion of the two particles in terms of the asymptotic traction solutions. In this way, the system remains fully coupled, incorporating both the boundary element representation for all other elements along with the near-field asymptotics.

As shown in Fig. 1, by combining theoretical analysis with hydrodynamic simulations, TC-BEM is able to correctly model both the far-field macroscale interactions and the near-contact microscale interactions. Using this new tool, we obtained the quantitative agreement of the self-diffusivities with the analytic model [5] for two rough spheres.




Fig. 1. TC-BEM simulated trajectories in the shear plane of two rough spheres suspended in a nonlinear shear flow.




Fig. 2. Migration diffusivity is a strong function of nonlinearity parameter.

The main result of the current study is that there is a net migration of the center of gravity of a pair of rough spheres suspended in a nonlinear shear flow towards the low-shear-rate region of the flow. This migration increases with the nonlinearity of the flow field and with particle roughness as shown in Fig. 2. This new theory was facilitated by our scale-coupling methodology and accurately predicted migration over a range of particle diameter from 0.1 to 3.175 mm.

**For more information contact Alan Graham at [graham@lanl.gov](mailto:graham@lanl.gov).**

- [1] D. Leighton, and A. Acrivos, *J. Fluid Mech.*, **181**, 415-439 (1987).
- [2] R.J. Phillips et al., *Phys. Fluids A*, **4**, 30-40 (1992).
- [3] G. Bossis, and J.F. Brady, *J. Chem. Phys.*, **87**, 5437-5448 (1987).
- [4] N. Tetlow et al., *J. Rheol.*, **42**, 307-327 (1998).
- [5] D.J. Jeffrey, and Y. Onishi, *ZAMP*, **35**, 634-641 (1984).

### Funding Acknowledgments

- Department of Energy
- Department of Energy, Office of Science, Office of Advanced Scientific Computing Research

## Simulations of Metallic Nuclear Fuel Casting


David Korzekwa, Deniece Korzekwa, MST-6; Neil Carlson, Bryan Lally, CCS-2; Mark Christon, Simulia

As part of the Global Nuclear Energy Partnership (GNEP), Idaho National Laboratory (INL) is designing a bench-scale casting furnace to develop the capability to cast thin rods of a uranium-plutonium-zirconium fuel alloy. The Telluride project [1] team and the MST-6 foundry at LANL are supporting this design effort by running simulations of casting processes to determine furnace design features and operability limits for the fuel casting process.

Fuel rods are long and thin, 380 mm long and 4.4 mm in diameter. The ultimate goal is to cast as many as 60 rods in a single pour. An initial bench-scale process is designed to cast three rods, as shown in Fig. 1. The multi-physics code Truchas is being used to simulate coupled fluid flow, heat transfer, and phase change, first for a single rod and subsequently for the three-rod bench-scale process.

The mold and crucible are made of graphite and heated by induction coils. Given the high thermal conductivity of the graphite, the mold must be filled very quickly to obtain a homogeneous casting. To achieve the needed fill rates, the mold will be evacuated and the metal injected under pressure. The simulation results show that for a constant pressure condition of 5 Kpa at the inlet (and nominal conditions elsewhere), the casting will freeze too quickly, preventing the mold from completely filling (Fig. 3). The simulations predict that the mold will fill successfully with an inlet pressure of 100 Kpa.

Future applications of Truchas to this problem will include operability studies to determine process parameters, sensitivity studies to quantify the accuracy required for pressure and temperature data, and improved material property models.



*Fig. 1. Geometry of the proposed bench scale mold and crucible.*




Fig. 2. Preliminary, single rod-filling simulation.




Fig. 3. Time sequence images of a rod-fill simulation. An inlet pressure of 5 KPa does not allow complete filling of the mold before solid metal obstructs the flow, which can be seen clearly in the last image. The radial direction is magnified 3x relative to the length to show more detail.

For further information contact David Korzekwa at [dak@lanl.gov](mailto:dak@lanl.gov).

[1] Telluride project: <http://telluride.lanl.gov>.

#### Funding Acknowledgments

- Department of Energy, Global Nuclear Energy Partnership
- Department of Energy, Office of Nuclear Energy



# A High-accuracy Limited-range Equation of State for Liquid Sodium

Eric Chisolm, T-1

A major focus of the equation of state (EOS) program in the past has been the production of EOS that cover very wide spans of compression and temperature and that have physically sensible behavior over the entire range, particularly the limits of extreme conditions. The SESAME library, maintained at LANL for many decades, epitomizes this approach. However, more modern applications such as the simulation of nuclear reactors require a different emphasis: the EOS need cover only a few thousand kelvin and a few atmospheres of pressure, but its predictions of thermodynamic properties must be exceedingly accurate in that range. We have made recent advances in the production of EOS to meet these needs, as illustrated by a new liquid sodium EOS that is valid up to roughly 1100 K and ten atmospheres and reproduces experimental data to a few percent. Such an EOS is required for the study of proposed reactors that use sodium as a coolant.

The goal was to find relations for energy  $E$  and pressure  $P$  as functions of density  $\rho$  and temperature  $T$  in the required regions. The available data consisted of the density  $\rho$ , constant-pressure specific heat  $C_p$ , thermal expansion coefficient  $\beta$ , and isothermal bulk modulus  $B_T$ , all as functions of temperature at a fixed pressure of one atmosphere. The EOS was constructed in two stages. First the energy at a pressure of one atmosphere was constructed by integrating the thermodynamic identity:

$$\left(\frac{\partial E}{\partial T}\right)_P = C_p - \frac{\beta P}{\rho}$$

which required all of the experimental data except for the bulk modulus. This determined the entire EOS at one atmosphere. To calculate the EOS off the isobar, we started with the expression:

$$P(\rho, T) = P^{iso}(T) + \frac{B_T^{iso}(T)}{\rho^{iso}(T)}(\rho - \rho^{iso}(T))$$

where the quantities with the superscript “iso” are given by the experimental data along the one atmosphere isobar. It is at this point that the isothermal bulk modulus is needed. This equation is not exact, but only an approximation valid at densities close to the isobar; it is here that we take advantage of the fact that the EOS is to be used only at low pressures. We integrated this expression to find the Helmholtz free energy  $F$ , and from there we calculated the energy  $E$ . The final expressions are complicated and may be found, with further supporting details, in [1]. Since we used an approximation for pressure, we needed to estimate its accuracy, concluding that the EOS is accurate to better than 5% in its least accurate form (and 0.5% in its most accurate form; see [1]) to pressures of 10 atmospheres.

We anticipate a greater demand for EOS of this type as more accurate simulations over smaller ranges of conditions become more important in nuclear energy and other applications.

**For further information contact Eric Chisolm at [echisolm@lanl.gov](mailto:echisolm@lanl.gov).**

[1] Eric Chisolm, “A high-accuracy limited-range equation of state for liquid sodium,” Los Alamos National Laboratory report LA-UR-06-6668.

## Funding Acknowledgments

- Department of Energy, Global Nuclear Energy Partnership






# Molecular Modeling of Alkaline Fuel Cell Membranes

Shaji Chempath, Lawrence R. Pratt, T-12; Brian R. Einsla, Clay S. Macomber, James M. Boncella, Jonathan A. Rau, Bryan S. Pivovar, MPA-11

Alkaline membrane fuel cells (AMFC) have the potential to replace the acidic proton-exchange membrane fuel cells (PEMFC) for use in vehicular and portable applications [1]. Alkaline media allow for the use of non-precious electrode catalysts made from inexpensive metals such as Fe/Co/Ni and thus avoid the use of rare precious metals such as Pt and Ru [2]. In AMFCs, electrolyte membranes that contain tetraalkylammonium head groups such as  $[\text{N}(\text{CH}_3)_4]^+$  are used. Unfortunately, hydroxide ions tend to react chemically with these head groups and form neutral species. This reaction between cation and hydroxide serves as the ultimate limit for membrane lifetime, and limits the applications where this technology can be used. We have used B3LYP/6-311++g(2d,p) level theory, polarizable continuum models, and reaction-path-finding algorithms [3] to study such reactions. We have also performed molecular dynamics simulations to probe the microstructure of alkaline membranes (Fig. 1).

Degradation of tetramethylammonium,  $[\text{N}(\text{CH}_3)_4]^+$ , is found to occur through two pathways as shown below (energies in kcal/mol). The reactant and product structures are also depicted in Fig. 2.



The small barrier (8.4 kcal/mol) calculated for ylide formation (Reaction 2a) suggests that rapid scrambling of methyl protons and aqueous protons should be observable in H-D isotope exchange experiments. Mass spectrometry of the gases [4] released in the thermal decomposition of  $[\text{N}(\text{CH}_3)_4]^+[\text{OD}]^- \cdot 5(\text{D}_2\text{O})$  does indeed reveal rapid scrambling of the protons of  $[\text{N}(\text{CH}_3)_4]^+$  with the  $[\text{OD}]^-$  and  $\text{D}_2\text{O}$ .

Transition state structures for different cations are shown in Fig. 3. In an alkaline membrane operating under fuel cell conditions, half of the volume is the non-aqueous phase (the polymer backbone) with the balance as the aqueous phase (water, cationic head groups, and  $\text{OH}^-$ ) as shown in Fig. 1. To probe the effect of water and solvation on reactivity we repeated the  $\Delta G^0$  and  $\Delta E^\ddagger$  evaluations with different dielectric constants in the PCM model and found that reactions are faster at lower water content (Fig. 4).

For further information contact Shaji Chempath at [shaji@lanl.gov](mailto:shaji@lanl.gov).

[1] J.R. Varcoe, C.T. Slade, *Fuel Cells*, **5**, 187 (2005).

[2] J.S. Spendlow, A. Wieckowski, *Phys. Chem. Chem. Phys.*, **9**, 2654 (2007).

[3] S. Chempath, C++ Source Code for Implementing the Growing String Method and Finding Transition States (<http://zeolites.cqe.northwestern.edu/shaji/growstring.html>).

[4] B.R. Einsla et al., *ECS Transactions*, **11**, 1173 (2007).

## Funding Acknowledgments

- Department of Energy, Office of Science, Office of Basic Energy Sciences




Fig. 1. Microstructure of alkaline membrane (blue spheres: fluorocarbon backbone, red: water).




Fig. 2. Two mechanisms for  $[N(CH_3)_4]^+$  degradation.




Fig. 3. Transition state structures for the  $S_N2$  attack on tetramethylammonium (TMA), phenyltrimethylammonium (phenyl-TMA), benzyltrimethylammonium (benzyl-TMA), and diazabicyclooctane (DABCO).




Fig. 4. The barrier ( $\Delta E^\ddagger$ ) for the degradation of  $[N(CH_3)_4]^+$  becomes small at low water content (at low  $\epsilon$ ).

# Accelerated Molecular Dynamics at the Solid-liquid Interface

Danny Perez, Arthur F. Voter, T-12

A long-standing problem in the atomistic simulation of materials is the fact that direct molecular dynamics (MD) techniques are limited to time scales of roughly one microsecond or less. Although computer speeds continue to increase, this limitation is not expected to change qualitatively in the near future. Parallel architectures can be used to increase the size scale of MD, but have little impact on the accessible time scale. For systems in which the long-time behavior can be characterized by infrequent events (typically thermally activated events) between well-defined states of the system, rate-based approaches can be powerful for overcoming this problem. In particular, accelerated molecular dynamics (AMD) methods developed over the last 10 years at LANL can often be used to access these longer time scales in an accurate, dynamical way. In the AMD approach, the core concept is to let the trajectory itself find an appropriate way to escape from each state, but in some way induce it to discover this escape path more quickly than it would under normal circumstances. In this approach, there is no need to guess at, or search for, all possible escape paths from each state, as in kinetic Monte Carlo (KMC) approaches. A discussion of both AMD and KMC methods can be found in [1].

There are many systems, however, that do not neatly fulfill the requirements for easy application of the AMD methods. A general class that is problematic is when the infrequent-event system is coupled with a complex and dynamic environment. The combined system makes much more rapid transitions, and direct application of AMD is pointless, since there is no computational gain. A specific example of this type of system is a solid surface in contact with a liquid. We might be interested in following the evolution of the surface morphology during electrochemical deposition or etching, or discovering the relevant steps of a surface-catalytic reaction or a corrosive process. As a step towards learning how to extend the AMD methods for treating this class of complex coupled systems, we have begun studying whether we can develop an AMD procedure that is valid for adatom diffusion at a model solid-liquid interface. This work also represents an important

step in the development, for an ongoing Scientific Discovery through Advanced Computing (SciDAC) project, of a multiscale simulation capability for stress-corrosion cracking.

As hinted above, AMD methods rely on a large separation of time scales between vibrations of atoms within a state and escape from the state. However, this separation does not hold in liquids where vibration and diffusion occur on similar time scales (on the order of a picosecond). A naive application of the AMD methods would thus be overwhelmed by the incessant change of state of the liquid, and thus could not provide acceleration of the dynamics of the solid atoms, which are really the quantity of interest here. The problem at hand thus amounts to defining the concept of a superstate so that the separation of time scales between intra- and intersuperstate dynamics is restored, and then applying AMD methods at the superstate level.

A simple strategy to reach this goal is to define superstates based on the configuration of the slow degrees of freedom (DOF) in the system, i.e., every state in which the slow DOF are in the same configuration belong to the same superstate, notwithstanding the state of the fast DOF. If the typical time scale over which the fast DOF evolve is much faster than that of the slow DOF, the validity of the AMD methods should be assured and their efficiency restored.

For typical solid-liquid interfaces—for example, a metal in contact with a liquid like water at room temperature—this segregation of DOF into fast and slow subsets should be easy to define: coordinates of the liquid atoms form the fast subset while coordinates of the solid atoms form the slow subset, while transitions in such a liquid will occur thousands of times faster than rearrangements of the solid atoms.

To test that this definition is proper, we carried out superstate parallel-replica dynamics [2] simulations of the diffusion of an adatom on a silver (100) surface (modeled using the Embedded Atom Method [3]) in contact with a film of a prototypical fluid (modeled using a Lennard-Jones potential). The distribution of 500 adatom hopping times at 600 K is shown in Fig. 1 for a direct simulation using conventional molecular dynamics and for a superstate parallel replica simulation on eight processors. The results clearly show that transition statistics are equivalent for the two approaches, implying that our superstate definition restores the time-scale separation essential for the validity of the AMD methods. In this case the parallel replica simulation was quite efficient, with a parallel efficiency of around 0.8 despite the presence of very fast transitions in the liquid.



## Materials

Note that the agreement is not due to a negligible effect of the liquid on the adatom dynamics. Indeed, as shown in Fig. 2, the transition rate of the adatom on the wet surface is around 1.6 times slower than on the dry surface at 600 K, increasing to a factor of 3.5 at 500 K. The rate corrections due to the liquid are thus significant and have a nontrivial temperature dependence. This further highlights the need to develop tools to study the peculiarities of the dynamics of solid-liquid interfaces.

In conclusion, we have demonstrated the possibility of adapting the AMD methods to systems in contact with fast-evolving complex environments. This ongoing work could help researchers tackle a wide range of technologically relevant problems like corrosion, liquid-metal embrittlement, or crystal growth and etching, by allowing these problems to be affordably studied at the atomic scale. While these results are encouraging, certain challenges remain, such as correctly treating the dynamics




Fig. 1. Distribution of hopping times for an adatom at a solid-liquid interface at 600 K.




Fig. 2. Temperature dependence of the hopping rate for an adatom at a dry interface and at a wet interface.

of dissolved solid atoms and generalizing our superstate definition to include lower temperature cases where the dynamics of the liquid in the solvation layer is not necessarily fast compared with the dynamics of the solid.

For further information contact Arthur Voter at [afv@lanl.gov](mailto:afv@lanl.gov).

- [1] A.F. Voter, F. Montalenti, and T.C. Germann, *Ann. Rev. Mater. Res.*, **32**, 321 (2002).
- [2] A.F. Voter, *Phys. Rev. B*, **57**, R13 985 (1998).
- [3] A.F. Voter, Los Alamos National Laboratory technical report LA-UR-93-3901 (1993).

### Funding Acknowledgments

- Department of Energy, Office of Science, Office of Advanced Scientific Computing Research
- Los Alamos National Laboratory Directed Research and Development Program

# Direct Transformation of Vacancy Voids to Stacking Fault Tetrahedral

Blas P. Uberuaga, Richard G. Hoagland, Steven M. Valone, MST-8;  
Arthur F. Voter, T-12

**I**t is well established that materials, upon irradiation, exhibit embrittlement and swelling that can often be traced to the atomic scale processes of defect accumulation. Defects (interstitials and vacancies) are produced; these diffuse and aggregate, leading to interstitial and vacancy clusters. When a supersaturation of vacancies is created, vacancy clusters can grow to the point that the material swells and fails. Here we report on our recent, unexpected observation that a vacancy void can thermally transform directly to a stacking fault tetrahedron (SFT) without passing through a Frank loop, the pathway previously postulated for the formation of SFTs in metals [1].

The simulations described here were done using parallel-replica dynamics (PRD) [2], an accelerated molecular dynamics method developed at LANL, in which time is parallelized. With PRD, given  $M$  processors, a speedup over molecular dynamics (MD) of up to a factor of  $M$  is possible. We used  $M=39$  and obtained speedups between 23 and 31. Nudged elastic band (NEB) calculations were then used to build state-to-state minimum energy paths. The Cu-Cu interaction was described using the embedded atom method (EAM) [3,4].

We performed PRD at constant volume on a 20-vacancy void at 400 K. After a phase lasting 1.35  $\mu\text{s}$  in which the extra vacancy diffused on the surface of the void, the extra vacancy reached an octahedral vertex, as illustrated in the inset of Fig. 1. Shortly thereafter, the void underwent a surprising transformation: as the vacancy executed a hop around the vertex, the corresponding moving Cu atom entered the void, triggering relatively direct collapse of the void into an SFT. Also surprising is that a process with such a high barrier (2.25–2.70 eV; see Fig. 1) could occur on this time scale. Assuming the rate is Arrhenius [ $k=\nu \exp(-\Delta E/k_B T)$ , where  $k_B$  is the Boltzmann constant] and a standard prefactor  $\nu=10^{13}/\text{s}$ , it would take  $\sim 10^6$  years for this event to occur. In contrast, we observe first-passage times  $\tau=1/k$  between 1 and 15 ns (once the void reaches the octahedral vertex state). Estimating the prefactor from the bold path in Fig. 1, with  $\tau=5.9$  ns and  $\Delta E=2.25$  eV, we find  $\nu=4 \times 10^{36}/\text{s}$ , a factor of  $10^{25}$  higher than the standard value. We also

find a Vineyard prefactor (computed from the normal modes of the system at the minimum and saddle [5]) of  $\nu=5 \times 10^{42}/\text{s}$ . These prefactors are more than 20 orders of magnitude larger than any prefactor we are aware of for processes in metallic systems, indicating that something unusual is happening.

This extremely high prefactor means that the transition state has a much higher entropy than the initial state, which we believe in this case results from an unusual volume effect. As the void collapses, the effective volume of the system increases as the material gains access to the empty space in the void—volume that was formerly excluded from it. For an isothermal system, when the system volume changes by  $\Delta V$ , the entropy changes by approximately [6]  $\Delta S=\alpha B \Delta V$ , where  $\alpha$  is the coefficient of volumetric thermal expansion and  $B$  is the bulk modulus. Using values for  $\alpha$  and  $B$  as predicted by this EAM potential, and assuming the transformation of a 20-vacancy void increases the effective system volume by 10 atomic volumes at the transition state, the entropy change is  $\Delta S=67.5 k_B$ . This predicts a rate prefactor increase of  $\exp(67.5) \sim 10^{29}$ , consistent with the observation above.

This void-to-SFT collapse can also occur for larger voids. For example, Fig. 2 shows the transformation of a 45-vacancy void, which collapsed after 0.25  $\mu\text{s}$  at  $T=475$  K. As in the case of the smaller void, the potential energy barrier is very large, over 4 eV, and the transformation is again driven by a large entropy increase.

Although it is not yet clear whether this void-collapse mechanism is important under typical radiation damage conditions, its existence illustrates that even the simplest of materials can behave in complex and unexpected ways that we are just beginning to understand. Accelerated molecular dynamics methods [7] offer a powerful tool for investigating these behaviors, as the dynamics often take place on time scales that are inaccessible to direct molecular dynamics.

**For further information contact A. F. Voter at [afv@lanl.gov](mailto:afv@lanl.gov).**

- [1] T. Jossang, J.P. Hirth, *Philos. Mag.*, **13**, 657 (1966).  
 [2] A.F. Voter, *Phys. Rev. B*, **57**, R13 985 (1998).  
 [3] M.S. Daw, M.I. Baskes, *Phys. Rev. B*, **29**, 6443 (1984).  
 [4] A.F. Voter, Los Alamos Technical report LA-UR-93-3901 (1993).  
 [5] G.H. Vineyard, *J. Phys. Chem. Solids*, **3**, 121 (1957).  
 [6] Y. Mishin, M. R. Sorensen, and A. F. Voter, *Philos. Mag. A*, **81**, 2591 (2001).  
 [7] B.P. Uberuaga, F. Montalenti, T.C. Germann, and A.F. Voter in *Handbook of Materials Modeling, Part A - Methods*, Ed. S. Yip, Springer, 629 (2005).

### Funding Acknowledgments

- Department of Energy, National Nuclear Security Administration
- Department of Energy, Office of Science, Office of Basic Energy Sciences




Fig 1. Inset: The 20-vacancy void in Cu. In this display convention, red spheres are vacancies. Notice the extra vacancy (labeled V) near one of the vertex sites (labeled C). Main Figure: Minimum energy path for the transformation of a 20-vacancy void to a SFT from 12 MD simulations started at the vertex state. The reaction coordinate has been shifted so that the highest energy saddle of each path is at 0. The circles and triangles represent the beginning and the end of each path.




Fig 2. Minimum energy path for the transformation of a 45-vacancy void to a SFT. The inset figures are minima from the parallel-replica simulation (red spheres are vacancies, and blue spheres interstitials, relative to an initial FCC reference structure). The open circles are estimates of the free energy along the path.

# Heterogeneous Nucleation and Hot Spots at Grain Boundaries in a RDX Bicrystal

Marc J. Cawkwell, Thomas D. Sewell, T-14; Kyle J. Ramos, Daniel E. Hooks, DE-9

A grain boundary is an interface between two crystals with different orientations. Grain boundaries are structural and elastic heterogeneities in a polycrystal and strongly influence a number of material properties [1,2]. Using non-equilibrium molecular dynamics (NEMD) simulations, we have studied the effects of a grain boundary on the response to shock loading of the energetic molecular crystal cyclotrimethylene trinitramine (RDX).

We constructed a bicrystal, with grains oriented as shown in Fig. 1, that measured  $39.7 \times 5.4 \times 100.0$  nm. Two parallel grain boundaries were created upon the application of periodic boundary conditions. The system was thermalized to 300 K and zero external pressure. The resulting grain boundaries displayed a period of about 13.8 nm over which the structure transitioned from semi-coherent to disordered. The bicrystal was impacted onto a stationary piston at a particle velocity of 500 m/s, generating a shock wave that propagated parallel to the plane of the grain boundaries.

Earlier NEMD studies of the propagation of shock waves parallel to [010] had revealed a shock-induced phase transformation [3]. These simulations employed a crystal free from pre-existing defects, hence the new phase nucleated homogeneously. Shock pressures around 10 GPa were required to nucleate even small amounts of the new phase. However, recent flyer-plate-driven shock wave experiments suggest that the phase transformation occurs at pressures of approximately 4 GPa. During our NEMD simulations of the bicrystal, we found that the new phase nucleates heterogeneously within the [010]-oriented grain from the disordered sections of the grain boundaries at a shock pressure around 4.5 GPa (Fig. 2). Hence, we attribute the discrepancy between our earlier theoretical predictions and recent experiments to the existence of heterogeneities that are inevitably present in bulk RDX single crystals.

We found no evidence for grain boundary sliding arising from the different shock propagation velocities in the two grains. However, we did measure a significant increase in intramolecular temperatures at the grain boundaries. This is shown in

Fig. 3 where we plot intramolecular temperatures calculated in slices parallel to the plane of the grain boundaries at the time of maximum compression during shock loading. The grain boundaries are 50-60 K hotter than the bulk crystals. Hence, grain or particle boundaries in energetic materials are hot spots that increase detonation sensitivity even in the absence of frictional heating due to interfacial sliding.

For further information contact Marc J. Cawkwell at [cawkwell@lanl.gov](mailto:cawkwell@lanl.gov).

- [1] A.P. Sutton, R.W. Balluffi, *Interfaces in Crystalline Materials*, Oxford University Press, Oxford (1995).
- [2] I. Alber et al., *Phil. Trans. R. Soc. Lond. A*, **339**, 555 (1992).
- [3] M.J. Cawkwell, T.D. Sewell, "New Shock-Induced Phase Transformation in RDX," Los Alamos National Laboratory report LALP-07-041 (2007).

## Funding Acknowledgments

- Joint memorandum of understanding (MOU) between the Department of Energy and the Department of Defense
- Department of Energy, National Nuclear Security Administration
- Los Alamos National Laboratory Directed Research and Development Program
- Office of Naval Research




Fig. 1. Schematic illustration of the bicrystal employed in the NEMD simulation.






Fig. 2. Section from the simulation cell during shock loading at a particle velocity of 500 m/s. The centers of mass of molecules identified as being in the  $\alpha$ -RDX environment are colored red while molecules belonging to the shock-induced phase are colored yellow. Arrows identify four heterogeneous nucleation events.




Fig. 3. Average intramolecular temperature calculated within 10-Å thick slices parallel to the grain boundary plane. The simulation cell was at the time of maximum compression in the NEMD simulation.



# Homogeneous Dislocation Nucleation and Anomalous Hardening in RDX

Marc J. Cawkwell, Thomas D. Sewell, T-14; Kyle J. Ramos, Daniel E. Hooks, DE-9

**T**he plastic deformation of crystalline solids is usually mediated by the motion of dislocations. Dislocations glide on crystallographic planes and the displacement they produce is known as the Burgers vector. Under typical loading conditions, dislocations nucleate and multiply heterogeneously by the formation and/or activation of dislocation sources. However, under shock loading, dislocation loops may nucleate homogeneously [1]. We recently discovered homogeneous nucleation of dislocations in the energetic molecular crystal cyclotrimethylene trinitramine (RDX) using non-equilibrium molecular dynamics (NEMD) simulations [2]. This unexpected deformation mechanism allowed us to quantitatively explain some puzzling experimental results.

Flyer-plate-driven shock waves with shock pressure  $P = 1.25$  GPa in (111)-oriented RDX single crystals revealed an elastic-plastic response that was expected based on the orientation of the shock with respect to the known slip systems [2]. However, at  $P = 2.25$  GPa we saw a single-wave, overdriven response [3] that suggested an abrupt change in the underlying mechanisms of plastic deformation with increasing pressure [2]. Interface velocity (VISAR) traces for shocks at  $P = 1.25$  and 2.25 GPa are presented in Fig. 1 (a) and (b), respectively.

Multi-million atom NEMD simulations of the propagation of planar shock waves normal to (111) were performed using an accurate and transferable potential for RDX. Crystal defects were not nucleated at  $P = 1.4$  GPa. However, for  $P \geq 1.8$  GPa, partial dislocation loops enclosing a stacking fault were nucleated homogeneously on (001). The Burgers vector of the partials was identified as  $\mathbf{b} = 0.16[010]$  by calculating the [010] cross-section of the (001)  $\gamma$ -surface [2,4]. The partial dislocations are highly mobile and rapidly generate stacking faults throughout the shocked material. Furthermore, analysis of the conformation of RDX molecules at the stacking fault showed that it is rendered metastable by a change in the orientation of the nitro groups with respect to mean plane of the six-member ring. In Fig. 2, we show only those molecules that have changed conformation, and hence the stacking faults, for a NEMD simulation at 2.2 GPa.

The stacking faults intersect the slip planes of the perfect dislocations that are responsible for the well-defined plastic wave seen experimentally at  $P = 1.25$  GPa. Thus, above the threshold shock pressure for homogeneous nucleation,  $1.4 < P_{th} \leq 1.8$  GPa, the material becomes plastically very hard because the stacking faults are obstacles to the glide of perfect dislocations. For this reason, above  $P_{th}$ , an overdriven rather than two-wave structure is found experimentally.

**For further information contact Marc J. Cawkwell at [cawkwell@lanl.gov](mailto:cawkwell@lanl.gov).**

- [1] B.L. Holian, P.S. Lomdahl, *Science*, **280**, 2085 (1998).
- [2] M.J. Cawkwell et al., *Phys. Rev. Lett.*, submitted (2008).
- [3] D.E. Hooks et al., *J. Appl. Phys.*, **100**, 024908 (2006).
- [4] V. Vitek, *Philos. Mag.*, **18**, 773 (1968).

## Funding Acknowledgments

- Joint memorandum of understanding (MOU) between the Department of Energy and the Department of Defense
- Department of Energy, National Nuclear Security Administration
- Los Alamos National Laboratory Directed Research and Development Program
- Office of Naval Research




Fig. 1. VISAR interface velocity traces for flyer-plate-driven shock waves in (111)-oriented RDX single crystals. (a)  $P = 1.25$  GPa, (b)  $P = 2.25$  GPa.



Fig. 2. NEMD simulation cell measuring 138.8 x 21.3 x 19.3 nm at the time of maximum compression during shock loading at  $P = 2.2$  GPa. The shock front propagated from left to right. The centers of mass of those molecules that changed conformation are shown. The stacking faults on (001) are visible as ribbons in this projection.

# Numerical Simulation of Projectile/Armor Interaction Using CartaBlanca

Xia Ma, Duan Z. Zhang, Paul T. Giguere, Qisu Zou, T-3

Accurate modeling and simulation of material strength and motion under ballistic or impact-loading conditions is a significant challenge and is of profound interest in this Laboratory. Recently, after significant improvements of the particle-in-cell (PIC) method and advancements of multimaterial interaction theory, we have developed a numerical method for such calculations.

On the modeling front, we developed a continuous multiphase flow or multimaterial deformation theory, and introduced a multipressure model. In this model each of the interacting materials possesses its own stress and pressure fields. A distinct advantage of this model is its ability to describe tension failure of a solid in coexistence with gas. In this model, the gas can remain under compression while allowing solid material to break under tension.

On the numerical front, a new algorithm is introduced to accurately satisfy the continuity constraint in a material point method (MPM) used for calculating multimaterial deformations. The MPM is an advanced version of PIC. Since its invention in the 1960s, the PIC method has been used in various problems, but the accuracy of satisfying the continuity constraint for multimaterials has limited its application. With the new algorithm, the continuity constraint is satisfied to a higher order of accuracy; therefore the MPM can be used to accurately calculate large deformations involving multimaterials and to avoid the numerical diffusion and mesh tangling issues associated with conventional Eulerian and Lagrangian methods.

Both the new theory and the new numerical method have found broad applications ranging from new design concepts of nuclear reactors, to mitigation of brain injuries due to improvised explosive device (IED) explosion, near-earth asteroid deflection, and the safety of liquid rocket fuels.

The figures shown in this article are from an armor-piercing calculation in which a cylindrical projectile is shot into high-hard armor steel. Figure 1 shows stress evolution in the steel during the penetration process. Our numerical results are compared with experimental data in Fig. 2.




Fig. 1. A tungsten rod penetrating a steel block (colors represent stress distribution).




Fig. 2. Comparison of numerical results with experimental data.

For further information contact Xia Ma at [xia@lanl.gov](mailto:xia@lanl.gov).

#### Funding Acknowledgments

- Department of Energy, Global Nuclear Energy Partnership

# Global-local Analysis of Laminated Composite Plates

Hashem M. Mourad, Todd O. Williams, Francis L. Addessio, T-3

Laminated composites are widely used structural materials, especially in aerospace, automotive, and defense applications. Their popularity can be attributed mainly to their high load-carrying capacity-to-weight ratio. Another characteristic of composite laminates is that their properties can be tailored—by changing, for example, the thickness, orientation, or stacking sequence of individual layers—with the goal of obtaining a desired response under given loading conditions.

To be able to optimize the properties of laminated composites in this manner, however, analytical and computational tools that can accurately and efficiently predict their behavior are needed. Conventional (equivalent single-layer) plate/shell theories do not resolve the layered setup and are therefore inadequate for modeling important damage mechanisms at the ply level, such as delamination. On the other hand, due to the small thickness of individual layers compared with the overall structural dimensions, general analysis techniques like 3D finite element methods are prohibitively expensive. Therefore, a new specialized analysis technique was developed.

The proposed technique [1] is a finite element formulation, based on the multiscale plate theory developed by Williams [2], and represents the first step toward the development of an efficient computational framework that can be used to study the response and damage distribution within composite laminates subjected to dynamic impact loading.

The formulation is built around the idea of expressing the displacement field as the summation of a) a global field, varying continuously over the thickness of the entire laminate, and b) a local field, whose definition varies from layer to layer. This layerwise definition of the local field allows discontinuities between layers (delamination) to be accounted for in a straightforward manner, and affords much improved accuracy, but incurs higher computational cost. However, the presence of a global component makes it possible to use the full global-local representation

in regions where high resolution is required (e.g., in the immediate vicinity of the collision site as in a large structure impacted by a small projectile), and a more economical global-only representation elsewhere, thereby maintaining overall computational efficiency.

Cohesive-zone models (CZMs) are used to predict the initiation and evolution of delamination. This enhances the predictive capability of the formulation, since CZMs do not require any assumptions about the pre-existence of delaminations in the structure. A hybrid finite element approach is adopted in which Lagrange multiplier fields, defined on the interface between adjacent layers, are used to enforce the appropriate interfacial constraint (the CZM or, in the perfectly bonded case, the persistent-contact condition).

Four-node (Q4T1) and nine-node (Q9T3) plate elements, based on this approach, were implemented and a set of standard benchmark problems were solved for verification purposes. One problem (Figs. 1 and 2) involves delamination in a simply supported cross-ply plate subject to pressure loading on its top surface. Another problem (Fig. 3), involving a cantilever plate subject to a transverse force, illustrates the robustness of the Q9T3 element in bending-dominated cases.




Fig. 1. Transverse displacement in a simply supported cross-ply laminate subject to pressure loading on its top surface. This square plate is relatively thick (with an aspect ratio of 10), and consists of three layers. Due to symmetry, only one quarter of the plate is modeled. The exact mid-point deflection is  $1.041 \times 10^{-3}$  in.






Fig. 2. Variation of the in-plane displacement,  $u_1$ , through the thickness of the laminate shown in Fig. 1 (at the mid-point of its edge coinciding with  $x_1=0$ ). The finite-element and analytical solutions are in excellent agreement. Interfacial displacement discontinuities due to delamination are clearly visible.

For further information contact Francis Addessio at [addessio@lanl.gov](mailto:addessio@lanl.gov).

[1] H.M. Mourad, T.O. Williams, and F.L. Addessio, "Finite element analysis of inelastic laminated plates: A global-local formulation with delamination," in preparation.

[2] T.O. Williams, *Int. J. Solids Struct.*, **36**, 3015–3050 (1999).

#### Funding Acknowledgments

- Joint memorandum of understanding (MOU) between the Department of Energy and the Department of Defense




Fig. 3. Transverse displacement in a cantilever plate subject to transverse loading at its tip. First-order elements (like Q4T1) are known to exhibit overly stiff behavior (locking) in such bending-dominated problems. The Q9T3 element was therefore used in this analysis and predicted the correct behavior. The exact tip deflection is 2.743 mm.

# Thermal Striping

David J. Torres, Rick M. Rauenzahn, Manjit S. Sahota, T-3

One facet of our Global Nuclear Energy Partnership (GNEP) activity has focused on modeling an experimental problem in which two liquid sodium jets at different temperatures impinge on a steel plate [1]. Thermal striping refers to the temperature fluctuations that are observed at the interface between two non-isothermal jets. This experiment is designed to reproduce a similar condition in the upper plenum of a liquid metal cooled fast breeder reactor (LMFBR). Heat is readily transferred to the steel, subjecting it to a repetitive cycle of temperature fluctuations that could potentially lead to fatigue and crack initiation.

We model the problem with the Computational Hydrodynamics for Advanced Design (CHAD) code, which is a parallel, finite-volume code capable of accommodating unstructured grids. Two inflow nozzles (one at 721 K and the other at 574 K) are placed 38 mm from a steel plate to match the experiment. The experiment is designed to be 2D by using nozzles with a rectangular cross section, thus we perform our calculations in a 2D grid. Figure 1 shows the dimensions of the computational domain. Each nozzle injects liquid sodium at 3.36 m/s into the computation. Two outflow boundaries at the upper left and upper right allow the sodium to exit. The steel is modeled using an isothermal equation of state, and the sodium is modeled with a Mie-Gruneisen equation of state. Calculations are run for an initial period of time to allow for the flow to develop.

Figure 2 shows a snapshot of the temperature just below the steel plate where the grid resolution is 1.4 mm in the horizontal direction and 0.475 mm in the vertical direction. The steel plate is physically above the region outlined by the box. The hot (red) sodium jet and the cooler (blue) sodium jet become unstable and interact to form eddies that travel outward from the center of the steel plate. One should note that sodium has a low kinematic viscosity so the Reynolds number is high ( $\sim 50,000$ ).

Figure 3 shows a time history of the temperature of sodium and that of steel at fixed points at the finest resolution (1.4 mm by 0.475 mm). The distances displayed in the legend are measured from the surface of the steel plate facing the

sodium jets. Thus (-82 mm, -2 mm) is -82 mm to the left of the horizontal middle of the steel surface and 2 mm down from the steel surface. The fluctuation is not a regular oscillation. The oscillation in the experiment is also not a regular periodic oscillation. The amplitude (and to some extent the frequency) of the oscillation in Fig. 3 depends on the location of the point sampled and ranges from approximately 3 K to 10 K at the surface of the steel to roughly 10 K to 45 K in the sodium. The amplitude quoted in the experiment is 20 K in the steel and 40 K in the sodium.

Fig. 4 shows the temperature oscillation in the sodium at three different grid resolutions at (-82 mm, -2 mm). The temperature in the lowest grid resolution (5.625 mm in the horizontal direction and 1.9 mm in the vertical direction) has been shifted down 40 K so that all three time-history plots at different resolutions could be compared. The frequency of oscillation at the lowest grid resolution can be identified and is 31 Hz. The behavior at the finer grid resolutions is more erratic,




Fig. 1. Computational dimensions.




Fig. 2. Sodium and steel temperature.

but still periodic in nature. The frequency quoted in the experiment is 14 Hz. However, if one counts the number of peaks in the experimental temperature history plot, the experimental frequency could be estimated to be between 23 and 28 Hz.

With CHAD, we have modeled an experiment relevant to reactors and demonstrated that the code can qualitatively reproduce the flow seen within the experiment and thus shed light on the complex physics encountered within reactors.

For further information contact David J. Torres at [dtorres@lanl.gov](mailto:dtorres@lanl.gov).

[1] K. Velusamy, et al., draft paper, *CFD Studies in the Prediction of Thermal Stripping in an LMFBF*.

#### Funding Acknowledgments

- Department of Energy, Global Nuclear Energy Partnership




Fig. 3. Temperature history in sodium and steel.




Fig. 4. Temperature history in sodium at three different grid resolutions.

# Perturbed Continuum Approach to Predict the Dynamic Behavior of Metals

Alek Zubelewicz, T-3; Jimmy Fung, X-2

Systems of various kinds that are subjected to non-equilibrium thermodynamic stimuli may incur irreversible damages leading to the breakdown of their functional capabilities. Often, the level of the damage is reduced and effectively counterbalanced through the redistribution of energy. We hypothesize that the energy rearrangement occurs along a well-defined thermodynamic path. At fixed positions of the path, the system maintains constant entropy. The positions are called the isentropic positions. An observer placed at the isentropic position may experience a rather smooth ride through the landscape of the thermodynamic system. However, the same ride monitored by an outside observer could appear as a quite violent event. Applicability of this hypothesis is tested for metals subjected to extreme loading rates. Our query is not so much whether a microstructural evolution occurs—we know that it does. Our interest is in answering a more fundamental question, namely what mechanisms necessitate the reconfiguration. For instance, a stressed metal develops a fairly regular array of dislocation cells, slip bands, and/or deformation twins. But it is not clear why the defects are so well organized at extreme loading rates. We investigate metal responses occurring within a very short time interval (in the range of  $10^{-7}$  to  $10^{-8}$  s) [1,2]. For this reason, a linear form of the strength model

$$\sigma_{eq} = (\sigma_0 + \theta_h e_{eq}^p) + \theta_\psi \psi_{inc}$$

seems acceptable. In this relation,  $\sigma_0$  is yield stress,  $\theta_h$  is plastic hardening defined at nanoscale (thus homogeneous at mesoscale), the fluctuating (plastic) strain is  $e_{eq}^p$ , and  $\theta_\psi$  is dynamic hardening due to the formation of mesoscale defect structures. The rate of plastic strain is coupled with the formation of incompatibilities  $\psi_{inc}$  in the field of particle velocity. Solution of the problem is obtained by satisfying the requirements of mass and momentum conservation. Expression for the dynamic overstress is found as a part of the solution.

A rather simple strength model with a built-in dynamic overstress is calibrated for copper and implemented into a 3D code. Several scenarios of the copper spherical

shell either imploding or expending under constant pressure are investigated. The analysis suggests that the stress-induced defects lead to a highly non-uniform pattern of deformation, which is in turn associated with a differential temperature. The imploding shell at a pressure of 1 GPa and at  $t=5.63$  ms is shown in Fig. 1. In the second case, an identical shell is subjected to a constant tensile pressure equal to 1 GPa. At an advanced stage of deformation ( $t=6.033$  ms), the stress-induced defects produce several hot spots leading to a breakup of the shell (Fig. 2).

**For further information contact Alek Zubelewicz at [alek@lanl.gov](mailto:alek@lanl.gov).**

[1] A. Zubelewicz et al., *Phys. Rev. B*, **71**, 104107 (2005).

[2] A. Zubelewicz, *Mechanics of Materials*, in press (2007).

## Funding Acknowledgments

- Department of Energy, Office of Science




Fig. 1. Imploding spherical shell made of copper subjected to a constant pressure (1 GPa) at  $t=5.63$  ms.




Fig. 2. Copper spherical shell at constant tensile pressure equal to 1 GPa and time  $t=6.033$  ms.



# Single Rod Fuel-clad Mechanical Interaction

Manjit S. Sahota, Rick M. Rauenzahn, David J. Torres, Thomas R. Canfield, T-3;  
Paul J. Maudlin, T-DO; Ralph. A. Nelson, ADWP

Much of our FY07 Global Nuclear Energy Partnership (GNEP) activity involved performing a number of coupled physics calculations in 1D and 3D using a single rod with fluid channel on the outside. Our particular focus was to simulate mechanical interaction between the fuel and cladding in case of severe fuel swelling that squeezes out the gap. We handled this interaction in the computational hydrodynamics for advanced design (CHAD) code by defining a critical value of the gap thickness (a user input parameter) below which the gap is considered to be closed. When this happens we add equal and opposite momentum sources at the two gap interfaces to maintain this thickness so that the cladding can respond to the fuel motion and vice-versa. (Simulating a flow blockage in the fluid channel is much easier as CHAD has built-in fluid velocity constraints at each node, which can be set to zero when the fluid channel thickness goes below a critical value.)

Figure 1 shows the fuel-clad interaction in 1D. CHAD is a 3D code, and all calculations are done in full 3D mode. To simulate a 1D rod, we took a small radial segment of the rod that has only one azimuthal zone and is only one cell high. Energy was added to the fuel while the pressure was applied to the outside surface of the coolant. We continued adding energy to the fuel until the gap was considered closed at 20 s. The fuel-clad interaction is depicted in the frames starting at 20 s.

Figures 2 and 3 show the fuel-clad interaction in 3D. Instead of the applied pressure to the coolant, the coolant is modeled as a flow channel with a solid cylindrical wall on the outside. The outside diameter of the flow channel is defined so that the total flow area corresponds to the advanced burner test reactor (ABTR) hexagonal design. Figure 3 compares the initial fuel-clad-coolant configuration with that at 25 s. An incipient Rayleigh-Taylor type instability is observable at 25 s. Soon after 25 s, the instability grows and the code fails to run further. The primary reason for this instability is that the deviatoric part of the material stress tensor was not included in these calculations, without which the fuel and the cladding behave like a fluid with very high density. The reason for not including the material strength

in these calculations is that the current semi-implicit implementation of these models in CHAD does not work satisfactorily for such low Mach number flows, which are of the order of about  $1 \times 10^{-8}$ . Work for fully implicitizing these models within the current implicit hydrodynamics scheme is in progress.

For further information contact Rick M. Rauenzahn at [rick@lanl.gov](mailto:rick@lanl.gov).

## Funding Acknowledgments

- Department of Energy, Office of Nuclear Energy



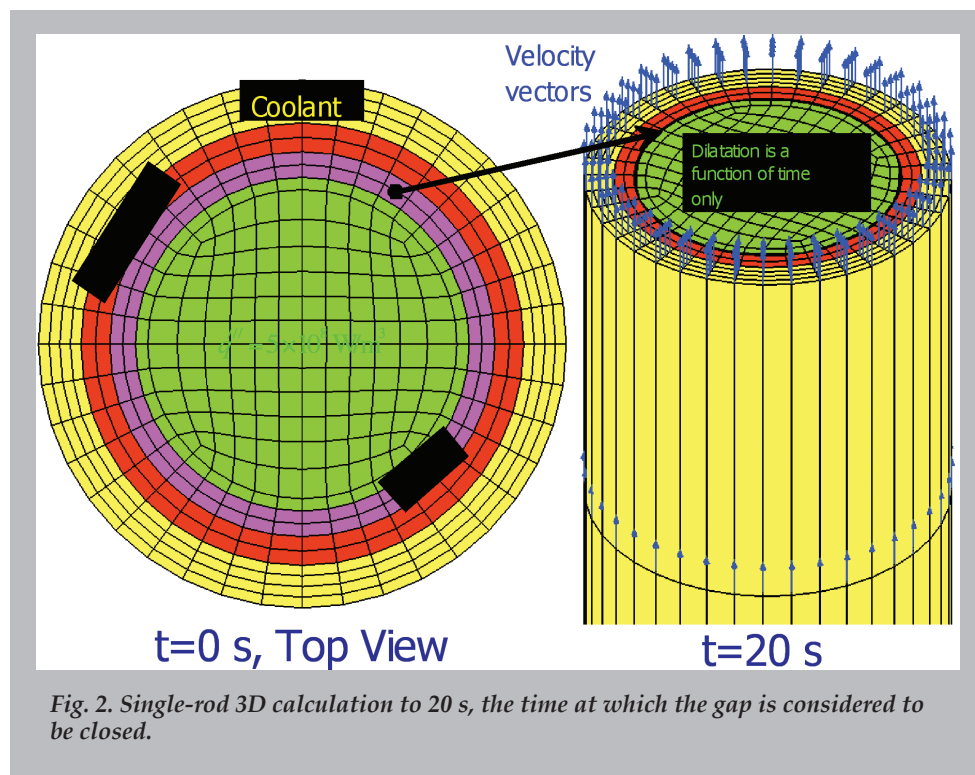


Fig. 1. Fuel-clad mechanical interaction in 1D cylindrical wedge.



# Disjoining Pressure for Non-uniform Thin Films

Bing Dai, Alan Graham, INST-OFF; Antonio Redondo, T-DO;  
Leslie G. Leal, Univ. of California, Santa Barbara

It is known that the concept of the disjoining pressure between two parallel interfaces due to van der Waals interactions cannot generally be extended to films of non-uniform thickness. We derive a formula for the disjoining pressure for a film of non-uniform thickness by minimizing the total Helmholtz free energy for a thin film residing on a solid substrate. The disjoining pressure for a thin film of small slope on a flat substrate is shown to take the form:

$$\Pi = -A_{123}(4 - 3h_x^2 + 3hh_{xx}) / 24\pi h^3$$

where  $A_{123}$  is the Hamaker constant for a phase 1 and 2 interacting through a phase 3;  $h$ ,  $h_x$ , and  $h_{xx}$  are the local film thickness, slope, and second order derivative, respectively.

Our work is concerned with a generalization of the theoretical approach, pioneered by Lifshitz and discussed more recently by Israelachvili, for the inclusion of van der Waals forces in a theoretical framework suitable for the dynamics of thin liquid films via the introduction of the so-called *disjoining pressure*.

In the continuum mechanics view, interfaces or boundaries between two macroscopic phases are treated as pure mathematical surfaces of zero thickness and zero mass. The macroscopic property of interfacial tension or interfacial free energy per unit area accounts for the difference in the collective molecular interaction effects for a heterogeneous system containing a fluid interface with respect to those interactions that would exist if the system were a homogeneous one. In a macroscopic multi-phase system, in which all phase boundaries and interfaces are sufficiently far apart, the incorporation of interfacial tension into the stress balance at fluid interfaces, together with the condition of continuity of velocities across the interface, produce a complete set of boundary conditions for free-surface problems. These conditions, together with the macroscopic Navier-Stokes equations, provide a basis for description of fluid motion and interface shapes. However, if the distance between interfaces or boundaries approaches

the mesoscopic regime ( $\sim 100$  nm or less), the collective effect of molecular interactions across these mesoscopic regions must be taken into account.

A rational revision is the *disjoining pressure* approximation, which was originally developed by Lifshitz for the case in which the mesoscopic region takes the form of a thin film. In this approach, which applies strictly for two unbounded parallel and flat interfaces or boundaries, the “additional” intermolecular forces between the phases across the film are replaced by an attractive force per unit area, which can be thought of as a pressure applied at the boundaries of the film, namely the so-called *disjoining pressure*:

$$\Pi = -A / 6\pi h^3$$

where  $h$  is the distance between the two parallel, sharp interfaces and  $A$  is the Hamaker constant, which can be evaluated via the Lifshitz theory. The simple concept of disjoining pressure between two parallel interfaces has been applied in a number of papers to study the stability and rupture of thin (but non-uniform) liquid films. However, as already noted, the simple formula for the disjoining pressure is exact only for thin films of uniform thickness. Moreover, as argued in some references, when the film thickness is very small,  $h \rightarrow 0$ , the classical formula for the disjoining pressure becomes unbounded. This leads to problems with the description of a thin film ending on a substrate and prevents determination of the motions near the edge of the film.

In our work, we have chosen to utilize a definition of the disjoining pressure based on Yeh’s thermal equilibrium condition for a liquid thin film residing on a flat solid substrate, but we modify it to include the excess energy outside the film. Assuming a van der Waals hard sphere interaction, we integrate the intermolecular potential throughout the whole heterogeneous system and obtain the excess energy associated with the disjoining pressure by subtracting the interfacial tension potential from the total intermolecular potential. In the limit of a small slope, the disjoining pressure given by this approach takes the relatively simple form:

$$\Pi = -A_{123}(4 - 3h_x^2 + 3hh_{xx}) / 24\pi h^3$$

where  $A_{123}$  is the Hamaker constant for the interactions between phase 1 and 3 through phase 2 and  $h$ ,  $h_x$ , and  $h_{xx}$  are the local film thickness, slope, and second-order derivative, respectively. A key point is that the form for the Hamaker constant depends on the properties of all three phases, and is completely consistent with Lifshitz theory for non-retarded van der Waals forces.

The extension to more general non-uniform thin films which are not necessarily bounded by a solid substrate is more or less trivial, and can be achieved by constructing tangential planes along the top and the corresponding bottom interfaces of the film.

One needs to be very careful when trying to extrapolate the concept of disjoining pressure, basically a mesoscopic concept, to molecular scales. The main problem arises from the characteristic length scales at which molecules are prevented from overlapping because of steric effects. Such steric effects automatically preclude extrapolation of the thickness of the film to zero values; they actually provide a minimum thickness which is of the order of molecular dimensions. On the other hand, in disjoining pressure analyses that start with a film thickness that is much larger than molecular length scales, the distinction between molecular scales and a film of zero thickness is usually ignored. But if the film thickness becomes comparable to molecular scales, the molecular size can no longer be neglected.

**For further information contact Alan Graham at [graham@lanl.gov](mailto:graham@lanl.gov).**

### **Funding Acknowledgments**

- Department of Energy, Office of Science, Office of Advanced Scientific Computing Research

# The Mechanism of Surfactant Effects on Drop Coalescence

Bing Dai, Alan Graham, INST-OFF; Antonio Redondo, T-DO;  
Leslie G. Leal, Univ. of California, Santa Barbara

**D**rop coalescence caused by collisions due to the flow of a suspending fluid is dramatically hindered by the presence of surfactants. The commonly accepted mechanism for this effect is that the fluid squeezed out of the contact region between adjacent drops generates a surfactant gradient. This in turn yields a Marangoni stress that is hypothesized to immobilize the interfaces and thus slow down the film drainage process that leads to coalescence. A scaling analysis based on an immobilized interface, however, is not consistent with recent experimental studies from Leal's laboratory at University of California, Santa Barbara (UCSB). Since the precise control of droplet size is of crucial importance in polymer blending and a number of other important technological processes, UCSB and LANL together initiated a multiscale mathematical investigation of the detailed coalescence mechanisms. The critical physics underlying this process spans length scales from macroscopic to molecular, as the film rupture process is driven by van der Waals forces across the thin film.

For our initial study, we attempt to simulate the axisymmetric head-on collision and coalescence of a pair of equal size Newtonian (though polymeric) drops in a Newtonian (though polymeric) suspending fluid, in the presence of an insoluble (polymeric) surfactant. For simplicity in the analysis, the viscosity of the suspending fluid is assumed to be the same as the viscosity of the drops. Interfacial dynamics for Stokes flow can then be represented by a boundary integral of the unknown surfactant concentration function and its derivatives

$$\mathbf{u}(\mathbf{x}_0) = \mathbf{u}^\infty(\mathbf{x}_0) - \frac{1}{8\pi Ca} \int \mathbf{G}(\mathbf{x}, \mathbf{x}_0) \cdot \left\{ \nabla_s \gamma(\Gamma) + [\Pi - \gamma(\Gamma)(\nabla \cdot \mathbf{n})] \mathbf{n} \right\} ds(\mathbf{x}) \quad (1)$$

here  $\mathbf{G}(\mathbf{x}, \mathbf{x}_0)$  is the free space Green's function while the whole term in the curly bracket is the stress difference along the interface. The stress difference is induced by the capillary pressure  $\gamma(\Gamma)(\nabla \cdot \mathbf{n})\mathbf{n}$ , the Marangoni stress  $\nabla_s \gamma(\Gamma)$ , and the disjoining pressure  $\Pi\mathbf{n}$  accounting for the van der Waals interactions at small scales. Where  $Ca$  is the capillary number, which measures the relative magnitude of viscous to capillary forces, and  $\mathbf{u}^\infty$  is the external flow. The drop shapes and positions are

updated based upon the kinematic condition using the velocity calculated from equation (1).

With the presence of surfactants, the interfacial dynamics is fully coupled with the surfactant transport, e.g., equation (2), that is solved via a Galerkin finite element method. The interfacial surfactant concentration is constructed from a set of Lagrangian quadratic basis functions defined on the moving mesh determined by the interfacial dynamics, e.g., equation (1).

$$\frac{\partial \Gamma}{\partial \tau} - \mathbf{U}_{\text{ref}} \cdot \nabla_s \Gamma + \nabla_s \cdot (\Gamma \mathbf{u}_s) - \frac{1}{Pe_s} \nabla_s^2 \Gamma + \Gamma(\mathbf{u} \cdot \mathbf{n}) \nabla \cdot \mathbf{n} = 0 \quad (2)$$

where  $Pe_s$  is the interfacial Péclet number, and the ratio of convective to diffusive flux and  $\mathbf{U}_{\text{ref}}$  accounts for the moving reference frame. A time-marching scheme is employed to exchange data between equations (1) and (2) for a convergence of both equations. Figure 1 shows the collision between two equal-sized drops with a uniform initial surfactant distribution in an axisymmetric, biaxial extensional flow. A surfactant gradient is built up along the interface, and the interfacial dynamics is completely different from that of the clean interface due to the Marangoni stress induced by the surfactant gradient.


Our benchmark results show that during coalescence, the interface in the thin gap actually exhibits a significant degree of mobility. This explains, at least to certain degree, the fact that the scaling behavior based on the assumption of a completely immobilized interface is not in agreement with recent experimental measurements. The unexpected new result is that the role of Marangoni effects on the coalescence process does not occur via immobilization of the interface within the thin gap region, but rather is due to its effects on the hydrodynamics outside the thin film. In particular, Marangoni stresses immobilize the drop interface outside the thin film, and this increases the total external hydrodynamic force that pushes the drops toward each other. This, in turn, increases the degree of flattening and dimpling of the thin film, and it is primarily this change that slows the film drainage process and thus increases the required drainage time prior to coalescence.


**For further information contact Alan Graham at [graham@lanl.gov](mailto:graham@lanl.gov).**

## Funding Acknowledgments

- Department of Energy, Office of Science, Office of Advanced Scientific Computing Research Program







300

350

400

5-month school clos



# Modeling Complex Networks

Few technical activities in the computational sciences have the same personal touch as do the modeling and simulation of complex networks. The term “complex network” refers to social, biological, or technological fabrics and interconnections with nontrivial topological structures. We are part of, interact with, and create new complex networks every day. Each time we interact with someone in our social network, search for information using Google, talk on a cellular telephone, drive to work, purchase an item, or consume electricity we are part of, and interact with, multiple active complex networks. The key to understanding such networks and extracting information from them is to discover within them islands of related entities and the connections between them.

The papers in this section demonstrate the many impacts of the application of knowledge in this field. They describe the modeling, simulation, and extraction of information from such networks—from the methods for and issues with discovering such islands and connectivities—to the extraction of actionable information from them that affect our lives and quality of life. For example, the study, comprehension, and modeling of complex networks can lead to understanding the impacts of and, more importantly, response to, a pandemic that will enhance life-saving measures and reduce the duration of such an event. Applications in this field can also be used to predict the impact of a natural disaster (or terrorist event) on a population and its infrastructure (and help plan mitigating measures), or can be used to understand and affect the social, environmental, and economic impacts of urbanization. This is an exciting technical field, at the threshold of discoveries that will have a positive effect on the lives of every human being on the planet.

# Structure Analysis of Large-scale Networks

Hristo Djidjev, CCS-3

One way to analyze and understand the structure of complex networks, and to find hidden relationships between their individual nodes, is to subdivide them into *communities*: maximal groups of related items. Computing the communities is of both theoretical and practical interest as they are used to identify groups of frequently interacting entities, such as collaboration groups, online societies, groups of scientific publications or news stories on a given topic, and related commercial items.

Finding the communities is one of a few, and often the only, available methods for analyzing the structure of a complex network. The importance of the problem has led to the development of a variety of algorithms based on different methods, yet no previous algorithm was able to offer *both* high accuracy and scalability. Our goal is to develop such an algorithm, and in particular, an algorithm that can compute community structures of networks of sizes of millions of nodes and edges, similar to the ones that appear in our simulations, in minutes or a few hours, and that will have accuracy comparable to the best available algorithms.

Although there is no widely accepted formal definition of a community decomposition of a network, intuitively it corresponds to a partition of the set of all nodes into subsets such that the density of *in-community links* (links between nodes of the same subset) is higher than the density of *between-community links*. Newman [1] formalizes this notion by defining a measure of clustering quality called *modularity*. We formulate the community detection problem as the problem of finding a partition maximizing the modularity, and we then solve the resulting optimization problem. The main features of our algorithm are as follows:

- We show that the problem of finding a partition maximizing the modularity of the original network is equivalent to solving a *minimum cut* problem on a modified network, which is a much better understood computer science problem.

- We solve the resulting minimum cut problem by applying a multilevel graph partitioning technique. In this method (Fig. 1), the original network is coarsened in a series of steps by merging pairs of adjacent nodes into single ones until a network of a sufficiently small size is reached. Then that small network is partitioned, and finally the partition of the smallest network is transformed into a partition of the original one by a series of projections and refinements.

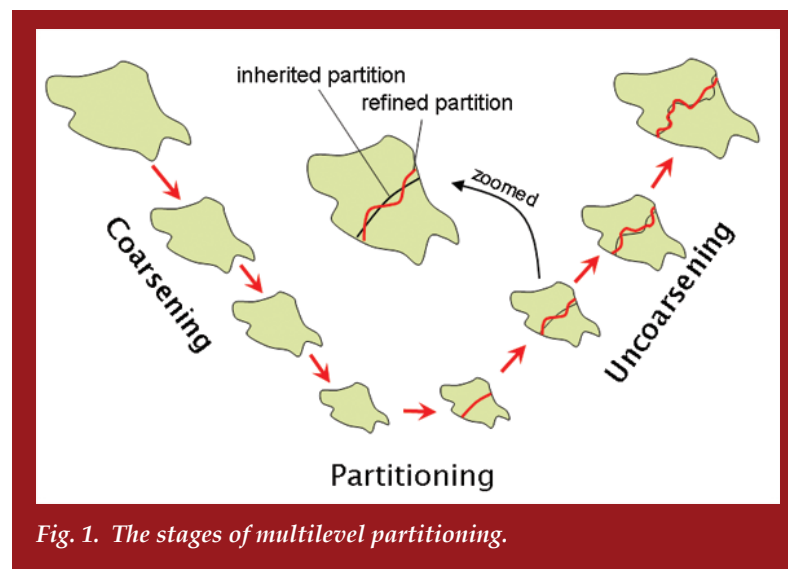


Fig. 1. The stages of multilevel partitioning.

Our experiments illustrate the high scalability and accuracy of our algorithm. We generated random networks of different sizes, densities, numbers of communities, and qualities of community structure. Figure 2 shows a comparison between the performance of our algorithm and that from [2], the only algorithm whose scalability comes close to our requirements. There are four groups of three experiments in each, comparing the performance of the algorithms with respect to the number of communities, the link densities, the strengths of communities, and power law parameters. As Fig. 2 shows, our algorithm is more accurate and about 30 times faster on average. Figure 3 shows that our algorithm is especially efficient for dense networks of small numbers of communities, analyzing a network of six million edges in about 15 seconds.



Nodes	Links	Com	$Q_{orig}$	$Q_{CNM}$	$Q_{ours}$	$Q_{=}$	$T_{CNM}$	$T_{ours}$
200	8930	2	.388	0	8	92	.61	.03
500	29801	5	.463	0	57	43	1.71	.23
900	71546	9	.387	1	99	0	4.44	.35
200	9932	2	.298	0	8	92	.68	.04
200	4967	2	.299	0	27	73	.54	.03
200	2458	2	.298	0	50	50	.61	.02
400	41856	4	.176	32	63	5	1.61	.18
400	43607	4	.154	39	60	1	1.66	.10
400	47797	4	.122	89	11	0	1.84	.07
400	8537	4	.244	0	100	0	1.35	.02
400	2653	4	.308	0	100	0	1.33	.03
400	888	4	.375	0	100	0	1.35	.02

Fig. 2. Comparison between the performances of our algorithm and the one from [2].  $Q_{orig}$  is the modularity of the partition used during network generation,  $Q_{CNM}$ ,  $Q_{ours}$ , and  $Q_{=}$  are the percentages of the cases where the algorithm from [2] produced a better modularity, our algorithm produced a better modularity, or both algorithms produced equal modularities, respectively.  $T_{CNM}$  and  $T_{ours}$  are the times of the algorithm from [2] and ours, respectively.

Nodes	Links	Total size	Time (sec)
5000	406125	411125	1.77
6000	764126	770126	3.09
7000	1283398	1300398	3.22
8000	1863710	1871710	4.66
9000	2418730	2427730	5.68
10000	3153106	3163106	7.27
15000	6295801	6310801	15.18

Fig. 3. Measuring the scalability of our algorithm.

For further information contact Hristo Djidjev at [djidjev@lanl.gov](mailto:djidjev@lanl.gov).

[1] M. Newman, *Phys. Rev. E*, **69**, 066133 (2004).

[2] A. Clauset, et al., *Phys. Rev. E*, **70**, 066111 (2004).

## Funding Acknowledgments

- Department of Homeland Security



## Using a Massive Agent-based Model to Study the Spread of Pandemic Influenza

Phillip D. Stroud, Sara Y. Del Valle, D-3; Susan M. Mniszewski, CCS-3;  
Jane M. Riese, Stephen J. Sydoriak, HPC-1

An influenza pandemic has great potential to cause large and rapid global and domestic increases in deaths and serious illness. Mathematical models and computer simulations have become the standard approach to studying epidemics. Recent analyses of nationwide pandemic dynamics [1,2] used detailed social contact models along with geospatial population density and travel patterns to develop agent-based models. However, they took average household size to be geospatially uniform, causing them to under-predict the variation and clustering of pandemic severity. Furthermore, they ignored the potential impact of temporary community-based strategies such as school closures.

We used the highly structured Epidemic Simulation System (EpiSimS) to simulate the detailed person-to-person contact of 19 million synthetic individuals representing the residents of Southern California. The EpiSimS agent-based discrete event engine [3] can simulate influenza pandemics at sufficient fidelity to capture geospatially varying demographic characteristics, travel patterns of individuals, and transmission opportunities through household, work, school, social, and casual contacts. We simulated a moderately severe influenza pandemic on a geographic stratum of 2.3 million locations, representing a six-county region of Southern California (Los Angeles, Orange, Riverside, San Bernardino, Ventura, and San Diego counties). Parameters describing the progression and transmission of influenza are based on [4]. Simulations were run on the Coyote cluster, an Institutional Computing resource.

The aggregate attack rate over the 3978 census tracts in the six counties is 30.6%, with a standard deviation of 7.5%. Figure 1 shows the attack rate for each census tract. The R-squared value of the *Pearson product moment correlation coefficient* between average household size and attack rate by tract is 0.90. Thus 90% of the variation in the attack rate by census tract is attributable to the average household size characteristic of the tract.

Results also demonstrate that temporary behavioral changes have the potential to generate waves, if they are relaxed before the pandemic dies out. The model

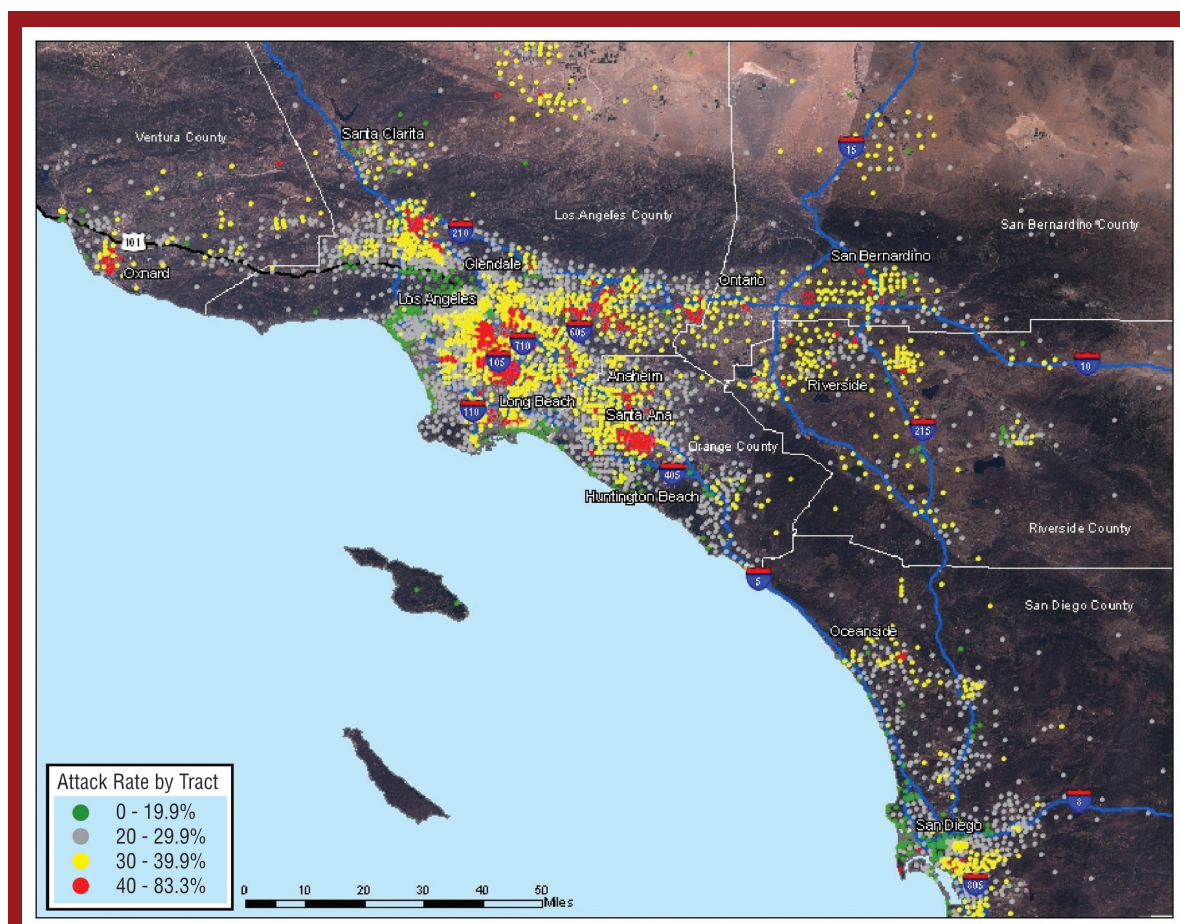


Fig. 1. The results of an individual-based stochastic model of the pandemic influenza in Southern California. The hot-spots are strongly correlated with regions of high average household size.

## Modeling Complex Networks

predicts that if school closures are implemented for the duration of the pandemic, the clinical attack rate could be reduced by more than 50%. Figure 2 shows the symptomatic percentage of the population as a function of time for the baseline and 100% school closures for 5, 8, and 11 months.

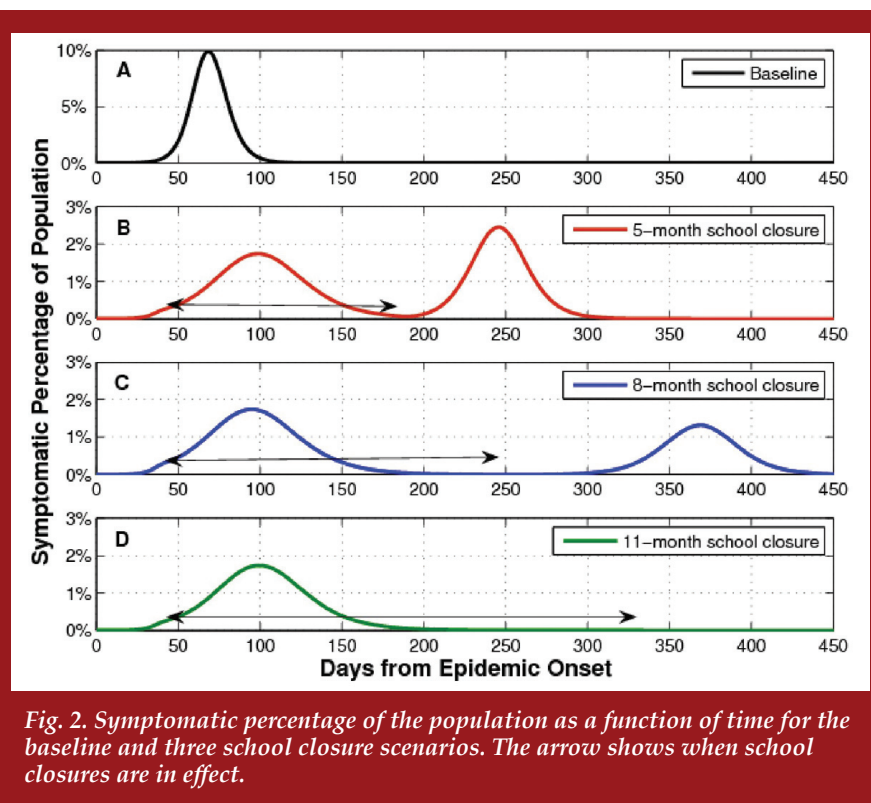
Our results are useful in providing guidance on the effects of school closures and deployment of limited resources. Readily available household-size statistics can be used to identify communities that are prone to higher-than-average attack rates. National or regional spatiotemporal pandemic simulation used to support influenza response planning should retain the geographic variation of demographic factors.

For further information contact Phillip D. Stroud at [stroud@lanl.gov](mailto:stroud@lanl.gov).

- [1] T.C. German, et al., *PNAS*, **103**(15), 5935-5940 (2006).
- [2] N.M. Ferguson, et al., *Nature*, **442**, 448-452 (2006).
- [3] P.D. Stroud, et al., *J Artif Soc Soc Simulat*, **10**(4) 9 (2007).
- [4] I.M. Longini, et al., *Am J Epidemiol*, **159**(7), 623-33 (2004).

### Funding Acknowledgments

- Department of Homeland Security



# Growth, Sustainability, Security, and the Pace of Life in Cities

Luís M. A. Bettencourt, T-7

**H**umanity has just crossed an historic threshold with a majority of people worldwide now living in cities. Yet, even as the debate on how humans impact the natural environment heats up, urbanization and its consequences remain poorly understood. For many, cities are the principal sources of our social and environmental problems in terms of crime and social unrest, pollution, poverty, and, often, incidence of several diseases. But cities have also always been disproportionately the birthplaces for most of human prosperity, innovation, and culture.

A team of researchers have found the key [1,2] to understanding and quantifying these seemingly contradictory features of urbanization. They analyzed a large number of urban indicators in the U.S., China, and several European countries, covering measures of economic productivity, innovation, demographics, crime, public health, infrastructure, and patterns of human behavior. They discovered that all these quantities follow simple statistical scaling relations (see Fig. 1) with population, indicating a continuum of change from small cities to the largest megapolis.

These relations quantify how material infrastructures can show savings per capita as people live more densely. But most interesting is the behavior of social and economic quantities: measures of wealth creation and innovation, among others, increase per capita with city size, in such a way that doubling the size of a city increases its economic productivity per person by about 15%. This is seen worldwide from China, to Europe, to the U.S.

What is fascinating and surprising about these results is that they show that the good things about cities—such as their innovation, and the bad ones—such as crime and the incidence of certain diseases, increase predictably in the same proportion as cities become larger. There is a continuum in all these quantities that accelerates dynamics that are already there in the smallest towns, making them apparent and conspicuous in the largest cities. This unity of dynamics,

characterized by faster and faster rates per capita with larger urban populations, means that the pace of life increases measurably with city size, as we have all experienced: cities are social accelerators. Having seen this effect in aggregated quantities such as wages, patents, or crime rates, the researchers could also see it in patterns of individual human behavior; for example, on average, people actually do walk faster the larger the city [1].

What these results mean for the growth of cities is also a fascinating issue. If growth is driven by economies of scale via exploitation of efficiencies, the city will ultimately stop growing. This has been known to economists for some time. But if instead, increasing returns to scale in innovation and wealth drive population growth, then a city will grow at an accelerated pace, faster than exponentially, and will eventually be driven to a crisis. To avoid this, cities must generate major adaptations that effectively restart these dynamics and begin a new growth cycle. The analysis also shows that the duration of such cycles decreases with population size. These patterns of accelerating, punctuated growth are observed for the historic population growth of New York City (see Fig. 1) and may be a feature of other social organizations.

These results will change the way people think of cities and project their growth. Although many of today's megacities present some of most daunting problems in terms of poverty, crime, and overtaxed infrastructure, these problems can be overcome by the ingenuity and resources that are already present in these places, as has happened historically in now-developed countries.

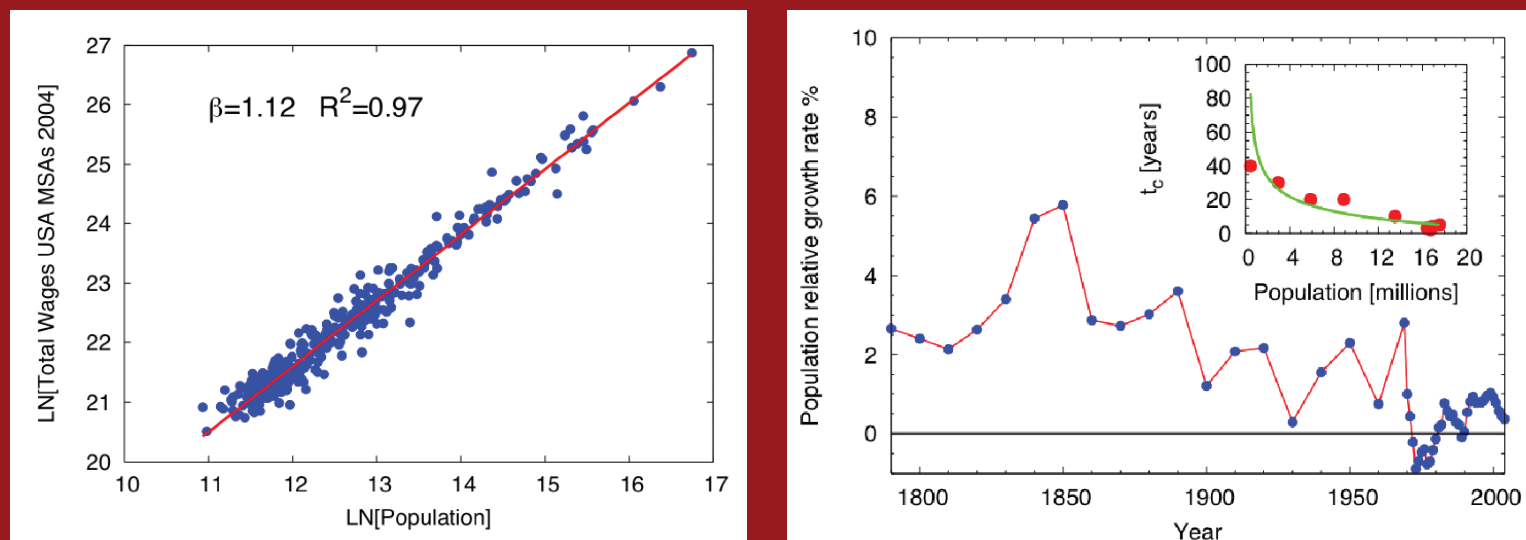
**For further information contact Luís M. A. Bettencourt at [lbett@lanl.gov](mailto:lbett@lanl.gov).**

[1] L.M.A. Bettencourt et al., *PNAS (USA)*, **104**, 7301-7306 (2007).

[2] L.M.A. Bettencourt et al., *Research Policy*, **36**, 107-120 (2007).

## Funding Acknowledgments

- European Union, Information Society as a Complex System Project
- Los Alamos National Laboratory Directed Research and Development Program
- National Science Foundation
- Thaw Charitable Trust



*Fig. 1. (Left) Increasing returns to scale in total wages for American Metropolitan Statistical Areas (MSA). Wages obey a scaling relation with exponent  $b=1.12>1$ , implying an increasing rate of wealth production as cities get larger. (Right) The population of New York City MSA has grown following shortening cycles of super-exponential growth. Inset shows the city's population growth as growth cycles shorten.*



## Predicting Mass Evacuation in the Wake of Natural or Human-caused Disasters

Leticia Cuéllar, Anders A. Hansson, Nick Hengartner, CCS-3; Deborah Kubicek, D-6

Suppose terrorists were to detonate radiological dispersion devices (RDDs), or “dirty bombs,” in Portland, Phoenix, and the U.S. Territory of Guam, causing widespread contamination. This was the scenario for *Top Officials 4*—the nation’s premier terrorism preparedness exercise—which took place in October 2007. To respond effectively, officials at the federal, state, territorial, and local level quickly needed information on numbers of casualties and affected individuals. This exercise and recent natural disasters show the need to develop contingency plans for emergency evacuation. In particular, rural communities and metro areas must be ready to absorb a large influx of evacuees, leading to increased demands on infrastructure and resources. It is, for example, estimated that one million people were evacuated due to Hurricane Katrina in 2005 (Fig. 1), of which at least 150,000 took refuge in Houston. Similarly, about half a million residents left their homes to seek shelter as the California wildfires raged earlier this year. To assist the Department of Homeland Security in its contingency planning effort, Leticia Cuéllar, Anders Hansson, and Nick Hengartner of CCS-3, together with Deborah Kubicek of D-6, have developed a fast-response tool to predict mass evacuation. The tool can guide pre-event resource allocation and aid emergency management teams and first responders to make timely and informed evacuation decisions.

Given a particular scenario/event, the tool generates a synthetic population of households for the area at risk. This is done in such a way that certain demographical characteristics (such as income, number of children, race) of the synthetic households are statistically equivalent to census data. We subsequently estimate the probability that each household will evacuate based on its demographics (e.g., number of children and elderly in the household, single-unit dwelling, mobile home) and risk data (actual and perceived risk, where the latter depends on previous experience). Similarly, risk and demographical data (income,


education, sex, and race) influence the household’s choice of shelter type: family/friends, hotel/motel, or public shelter. The tool finally makes use of a gravity model in which the attractiveness of a shelter increases with the number of additional evacuees it can harbor, but declines with increasing remoteness, in order to geo-locate the preferred shelter type.

The output of the tool informs the decision maker about a) the influx of evacuees to neighboring communities (expressed in number of people, broken down by shelter type preference, in Figs. 1,2), and b) the number of people that neglect to evacuate in each impacted census block group. These results are readily analyzed and visualized using a geographic information system. The output can also be shared with other analysis tools employed by LANL’s National Infrastructure Simulation and Analysis Center (NISAC) to assess the impact on critical infrastructure and key resources.

**For more information contact Anders Hansson at [hansson@lanl.gov](mailto:hansson@lanl.gov).**


### Funding Acknowledgments

- Department of Homeland Security

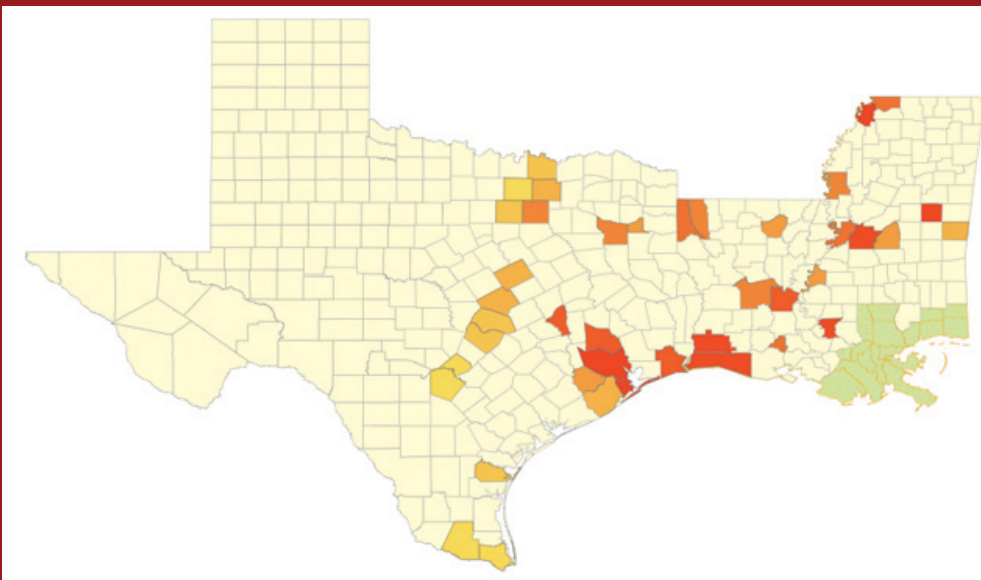


*Fig. 1. Hurricane Katrina near peak strength on August 28, 2005.*





*Fig. 2. Predicted influx of evacuees to neighboring family/friends in the wake of Hurricane Katrina. Darker red indicates a higher influx.*



*Fig. 3. Predicted influx of evacuees to neighboring hotels/motels in the wake of Hurricane Katrina. Darker red indicates a higher influx.*

## Evaluation of First Born Program for Los Alamos National Laboratory (LANL) Foundation

Aparna V. Huzurbazar, CCS-6

Children in rural New Mexico are more likely to experience adverse life events such as child abuse or parental substance abuse in childhood, and youth and adults in this region exhibit extremely poor outcomes relative to their peers around the U.S. To help address this problem, a public/private partnership, which includes the LANL Foundation and the State of New Mexico, is implementing the First Born Program (FBP) in Rio Arriba and Taos Counties. FBP is an evidence-based, home visitation program that uses clinically trained professionals to provide parents and families with the information, training, and resource access, promote the highest level of development for a newborn infant up to age three—the most important period for early brain development.

The logic model of this intervention is a reasonable one, and longitudinal studies of other child development programs indicate that early and continuous intervention can have a positive impact on a person's early childhood and lifelong development. However, a sound logic model and supporting evidence from other interventions is not a guarantee that a prevention program will have its intended effects. This study undertakes an impact evaluation to assess the effectiveness of this model in New Mexico.

**Research Design.** The Northern New Mexico implementation of FBP is a community-based universal model. The research design most applicable to this program setting is a pretreatment cohort control group design. This design involves comparing outcomes of eligible families after the FBP began operating to outcomes of target families who would have been eligible for FBP but gave birth before the program commenced. Data collection will begin this spring. We plan to collect data that will enable us to estimate a regression model that identifies the treatment and other factors that explain the ultimate and intermediate outcomes for both children and mothers. Our objective is to collect data until study children are eight years old.

**The Evaluation Team.** LANL statistician Aparna V. Huzurbazar will lead the project through an appointment at the New Mexico Consortium. The RAND Corporation of Santa Monica, CA, is leading the study with senior economist Rebecca Kilburn as the Project Leader. The project takes advantage of partnerships and assets in New Mexico, and peer review from nationally recognized leaders in the field.

**For further information contact Aparna V. Huzurbazar at [aparna@lanl.gov](mailto:aparna@lanl.gov).**

### Funding Acknowledgments

- Con Alma Foundation
- Los Alamos National Laboratory Foundation



*Fig. 1. Pictured are family Isidor Chavez, Marie McClard, and daughter Izabella (Bella) Chavez. Photo by Don Usner, LANL Foundation.*



